

Microsoft.AZ-104J.v2026-10-22.q188

試験コード :

AZ-104J

試験名称 :

Microsoft Azure Administrator (AZ-104日本語版)

認証ベンダー :

Microsoft

無料問題の数 :

188

バージョン :

v2026-10-22

ページの閲覧量 :

115


問題集の閲覧量 :

2282

<https://www.jpnsshiken.com/shiken/Microsoft.AZ-104J.v2026-10-22.q188.html>

質問: 1

次の図に示すストレージ アカウントを含む Azure サブスクリプションがあります。



ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

You can create a premium file share in

contoso101only

contoso104 only

contoso101 or contoso104 only

contoso101, contoso102, or contoso104 only

contoso101, contoso102, contoso103, or contoso104

You can use the Archive access tier in

contoso101only

contoso101 or contoso103 only

contoso101, contoso102, and contoso103 only

contoso101, contoso102, and contoso104 only

contoso101, contoso102, contoso103, and contoso104

正解:

You can create a premium file share in

contoso101only

contoso104 only

contoso101 or contoso104 only

contoso101, contoso102, or contoso104 only

contoso101, contoso102, contoso103, or contoso104

You can use the Archive access tier in

contoso101only

contoso101 or contoso103 only

contoso101, contoso102, and contoso103 only

contoso101, contoso102, and contoso104 only

contoso101, contoso102, contoso103, and contoso104

Explanation:

You can create a premium file share in

contoso101only

contoso104 only

contoso101 or contoso104 only

contoso101, contoso102, or contoso104 only

contoso101, contoso102, contoso103, or contoso104

You can use the Archive access tier in

contoso101only

contoso101 or contoso103 only

contoso101, contoso102, and contoso103 only

contoso101, contoso102, and contoso104 only

contoso101, contoso102, contoso103, and contoso104

In Microsoft Azure Storage, the type and kind of a storage account determine which features and performance tiers are supported.

1## Premium File Share Support:

Premium file shares are only supported in FileStorage accounts, which provide high-performance, low-latency SSD-based storage optimized for enterprise file workloads. According to the official Azure Storage documentation (Microsoft Learn: " Azure Files performance tiers " ), only the FileStorage account kind supports Premium file shares.

In the exhibit, only contoso104 is a FileStorage account. Therefore, premium file shares can only be created in contoso104.

2## Archive Access Tier Support:

The Archive access tier is available for Blob storage and General-purpose v2 (StorageV2) accounts that store block blobs. The Archive tier is used for data that is rarely accessed and provides the lowest storage cost.

From the exhibit:

- \* contoso101 (StorageV2) # supports archive tier.
- \* contoso102 (Storage) # older type, does not support access tiers beyond hot/cool.
- \* contoso103 (BlobStorage) # supports archive tier.
- \* contoso104 (FileStorage) # does not support Blob tiers.

Therefore, only contoso101, contoso102, and contoso103 allow the Archive tier, but since standard "Storage" (contoso102) doesn't support Archive, the correct accounts are contoso101 and contoso103.

質問: 2

次の表に示すユーザーを含む Azure AD のハイブリッド展開があります。

| Name  | User type | On-premises sync enabled |
|-------|-----------|--------------------------|
| User1 | Member    | No                       |
| User2 | Member    | Yes                      |
| User3 | Guest     | No                       |

ユーザーの JobTitle 属性と UsageLocation 属性を変更する必要があります。

Azure AD から属性を変更できるのはどのユーザーですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Microsoft

JobTitle: User1 and User3 only

User1 only

User1 and User2 only

User1 and User3 only

User1, User2, and User3

UsageLocation: User1, User2, and User3

User1 only

User1 and User2 only

User1 and User3 only

User1, User2, and User3

正解:

Answer Area

Microsoft

JobTitle: User1 and User3 only

User1 only

User1 and User2 only

User1 and User3 only

User1, User2, and User3

UsageLocation: User1, User2, and User3

User1 only

User1 and User2 only

User1 and User3 only

User1, User2, and User3

Explanation:

Answer Area

JobTitle: User1 and User3 only

UsageLocation: User1, User2, and User3

In a hybrid Azure Active Directory (Azure AD) environment, where user identities can be cloud-only or synchronized from on-premises Active Directory (AD), attribute management depends on the source of authority for each user account.

The source of authority determines where you can modify a user's attributes:

If a user is synced from on-premises Active Directory (on-premises sync enabled = Yes), certain attributes are read-only in Azure AD and must be edited in the on-premises AD. These attributes include:

- Job title
- Department
- Office
- Manager
- Company name

If a user is cloud-only (not synchronized) or a guest user, you can modify those same attributes directly in Azure AD via the Azure portal, Microsoft Graph, or PowerShell.

According to Microsoft Learn (Azure AD attributes and synchronization rules):

"When Azure AD Connect synchronizes identities, certain user attributes such as job title, department, and manager are mastered in on-premises Active Directory and become read-only in Azure AD. However, attributes like usage location can always be edited in Azure AD, as they are required for license assignment." Applying to this question User Type On-premises sync enabled Editable in Azure AD User1 Member No Editable (Cloud-only) User2 Member Yes On-premises sync (Read-only for JobTitle) User3 Guest No Editable in Azure AD JobTitle User1 - Editable (cloud-only member) User2 - Not editable (synced from on-prem AD) User3 - Editable (guest accounts managed in Azure AD)

# Answer: User1 and User3 only

UsageLocation

This attribute must be set in Azure AD for licensing and service access.

It can be edited for all users - even for synced users or guests - because Azure AD maintains it as a cloud- only attribute, not synchronized from on-prem AD.

# Answer: User1, User2, and User3

Final Verified Answers (Based on Microsoft Azure Administrator Documentation):

Attribute

Editable Users

JobTitle

# User1 and User3 only

UsageLocation

# User1, User2, and User3

Microsoft Official Documentation Extract:

"Attributes such as job title, department, and office location are sourced from on-premises AD and cannot be modified in Azure AD for synchronized users. The usage location attribute, however, is always managed in Azure AD." (Source: Microsoft Learn - Azure AD Connect sync: Attributes synchronized to Azure AD and Manage user profile attributes in Azure AD)

質問: 3

Subscription1 という名前の Azure サブスクリプションがあります。Subscription1 には、次の表に示すリソース グループが含まれています。

| Name | Azure region   | Assigned Azure Policy |
|------|----------------|-----------------------|
| RG1  | West Europe    | Policy1               |
| RG2  | North Europe   | Policy2               |
| RG3  | France Central | Policy3               |

RG1にはWebApp1というWebアプリがあります。WebApp1は西ヨーロッパにあります。

WebApp1 を RG2 に移動します。

この移転による影響は何ですか？

- A. WebApp1 の App Service プランが北ヨーロッパに移動します。WebApp1 にはポリシー 2 が適用されます。
- B. WebApp1 の App Service プランは西ヨーロッパのままです。ポリシー 2 が WebApp1 に適用されます。
- C. WebApp1 の App Service プランは西ヨーロッパのままです。ポリシー 1 が WebApp1 に適用されます。
- D. WebApp1 の App Service プランが北ヨーロッパに移動します。Policy1 が WebApp1 に適用されます。

正解: [\(正解を表示します\)](#)

質問: 4

App1とApp2という名前の2つのAzure App Service Webアプリがあります。それぞれのWebアプリには、運用環境へのデプロイ用スロットとテスト環境へのデプロイ用スロットがあります。

ウェブアプリの運用スロットにおけるバックアップ構成設定を以下の表に示します。

| Web app name | Backup Every | Start backup schedule from | Retention (Days) | Keep at least one backup |
|--------------|--------------|----------------------------|------------------|--------------------------|
| App1         | 1 Days       | January 6, 2024            | 0                | Yes                      |
| App2         | 1 Days       | January 6, 2024            | 30               | Yes                      |

以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注：正解ごとに1ポイントが加算されます。

Answer Area

Statements

Microsoft

On January 15, 2024, App1 will have only one backup in storage.

On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.

On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.

Yes

No

☐

☐

☐

☐

☐

☐

正解:

Answer Area

Statements

Microsoft

On January 15, 2024, App1 will have only one backup in storage.

On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.

On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.

Yes

No

☒

☐

☒

☐

☒

☐

Explanation:

Answer Area

Statements

Microsoft

On January 15, 2024, App1 will have only one backup in storage.

On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.

On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.

Yes

No

☒

☐

☒

☐

☒

☐

質問: 5

share1 という名前のファイル共有を含む、storage1 という名前の Azure ストレージ アカウントを作成する予定です。

share1 が SMB マルチチャネルをサポートしていることを確認する必要があります。ソリューションはコストを最小限に抑える必要があります。

ストレージ1をどのように構成すればよいですか？

- A. ローカル冗長ストレージ (IRS) を使用した標準パフォーマンス
- B. ローカル冗長ストレージ (LRS) によるプレミアム パフォーマンス

**C. ゾーン冗長ストレージ (ZRS) による標準パフォーマンス**

正解: [\(正解を表示します\)](#)

SMB Multichannel is a feature of Azure Files that enables clients to establish multiple network connections to an Azure file share simultaneously. This allows parallel data transfer, thereby increasing throughput and providing fault tolerance for high-performance workloads.

According to the Microsoft Azure Administrator documentation and Azure Files technical reference, SMB Multichannel is available only for premium file shares hosted on Azure Files Premium Storage accounts that use locally redundant storage (LRS) replication.

Extract from Microsoft documentation:

"SMB Multichannel is supported only for premium file shares. Premium file shares are hosted in FileStorage storage accounts, which use SSD-based performance and locally redundant storage (LRS). Standard (HDD- based) file shares do not support SMB Multichannel." The Premium tier uses SSD storage, designed for low latency and high IOPS workloads, and is required for enabling SMB Multichannel. Using LRS (Locally Redundant Storage) provides the lowest cost redundancy option while still ensuring three replicas within the same datacenter.

Therefore, to meet the requirement of SMB Multichannel while minimizing costs, you must choose Premium performance with LRS redundancy.

**質問: 6**

Azure サブスクリプションをお持ちです。このサブスクリプションには、VNet1 という仮想ネットワークに接続する仮想マシンが含まれています。

VM Insights 用に Azure Monitor を構成する予定です。

すべての仮想マシンが VNet1 経由でのみ Azure Monitor と通信するようにする必要があります。

最初に何を作成すればよいですか？

**A.** Azure Monitor プライベート リンク スコープ (AMPLS)

**B.** プライベートエンドポイント

**C.** Log Analytics ワークスペース

**D.** データ収集ルール (DCR)

正解: [\(正解を表示します\)](#)

Azure Monitor for VM Insights is a feature of Azure Monitor that provides comprehensive monitoring and diagnostics for your Azure virtual machines and virtual machine scale sets. It collects performance data, process information, and network dependencies from your virtual machines and displays them in interactive charts and maps. You can use Azure Monitor for VM Insights to troubleshoot performance issues, optimize resource utilization, and identify network bottlenecks1.

To enable Azure Monitor for VM Insights, you need to install two agents on your virtual machines: the Azure Monitor agent (preview) and the Dependency agent. The Azure Monitor agent collects performance metrics and sends them to a Log Analytics workspace. The Dependency agent collects process information and network dependencies and sends them to the InsightsMetrics table in the same workspace2.

By default, the agents communicate with Azure Monitor over the public internet. However, if you want to ensure that all the virtual machines only communicate with Azure Monitor through a virtual network named VNet1, you need to configure private network access for the agents.

Private network access allows the agents to communicate with Azure Monitor using a private endpoint, which is a special network interface that connects your virtual network to an Azure service without exposing it to the public internet. A private endpoint uses a private IP address from your virtual network address space, so you can secure and control the network traffic between your virtual machines and Azure Monitor3.

To configure private network access for the agents, you need to create an Azure Monitor Private Link Scope (AMPIS) first. An AMPIS is a resource that groups one or more Log Analytics workspaces together and associates them with a private endpoint. An AMPIS allows you to manage the private connectivity settings for multiple workspaces in one place4.

After creating an AMPIS, you need to create a private endpoint in VNet1 and link it to the AMPIS. This will enable the agents on your virtual machines to send data to the Log Analytics workspaces in the AMPIS using the private IP address of the private endpoint5.

**質問: 7**

ネットワークにはオンプレミスの Active Directory ドメイン サービス (AD DS) ドメインが含まれています。

ドメインには、次の表に示す ID が含まれています。

| Name   | Description                      | In organizational unit (OU) |
|--------|----------------------------------|-----------------------------|
| User1  | User                             | OU2                         |
| User2  | User                             | OU1                         |
| Group1 | Global group that contains User1 | OU1                         |

Microsoft Entra Connect 同期を使用して OU1 を AD DS ドメインから Microsoft Entra テナントに同期する Azure サブスクリプションがあります。  
Microsoft Entra テナントには、User3 という名前のクラウドのみのユーザーが含まれています。

Answer Area

Statements

User1 can access content in share1.

User2 can access content in share1.

User3 can access content in share1.

Yes

No

正解:

Answer Area

Statements

User1 can access content in share1.

User2 can access content in share1.

User3 can access content in share1.

Yes

No

Explanation:  
NO  
YES  
NO

In this scenario, Microsoft Entra Connect (formerly Azure AD Connect) is configured to synchronize only the Organizational Unit (OU1) from the on-premises Active Directory Domain Services (AD DS) environment to the Microsoft Entra tenant (Azure AD).

The synchronization scope determines which objects (users, groups, devices) are replicated to Azure AD. Any objects outside the selected OUs are not synchronized and therefore do not exist in Azure AD.

Here's what happens to each user:

- \* User1 (in OU2) - Because OU2 is not included in the synchronization scope, User1 is not synchronized to Azure AD. Therefore, User1 does not have an account in Microsoft Entra ID and cannot authenticate to any cloud-based resource such as Share1 (which is assumed to be an Azure file share or Microsoft 365 resource).
- \* Result: No access.
- \* User2 (in OU1) - Since OU1 is synchronized using Microsoft Entra Connect, User2 exists in Azure AD. User2's identity is recognized by Azure AD and can be used to access cloud-based resources (like Share1) if permissions are assigned appropriately (either directly or via a group such as Group1).
- \* Result: Yes, access is possible.
- \* User3 (cloud-only account) - While User3 exists in Azure AD, there's no indication that this user has been assigned access to Share1. Because access is granted through synchronization (from AD DS) and group membership in synchronized objects (like Group1), User3 would not automatically inherit those permissions.
- \* Result: No access.

Additionally, Group1, which contains User1 and resides in OU1, will be synchronized. However, since User1 is in OU2, User1's membership in Group1 is not synchronized, because non-synchronized objects cannot appear in synchronized group memberships in Azure AD.

This behavior is defined by Microsoft's Azure AD Connect synchronization rules, which state:

"Objects outside of the configured synchronization scope are not synchronized to Azure AD. Group memberships including users from outside the synchronization scope will not include those users in Azure AD."

質問: 8

お客様は、contoso2023 という名前のストレージ アカウントを含む Azure サブスクリプションをお持ちです。このストレージ アカウントには、以下のものが含まれています。

\* cont1: ブロブコンテナ

\* share1: ファイル共有



Performance ⓘ

☒ Standard ☐ Premium

**i** This setting cannot be changed after the storage account is created.

Secure transfer required ⓘ

☐ Disabled ☒ Enabled

Allow Blob anonymous access ⓘ

☐ Disabled ☒ Enabled

Allow storage account key access ⓘ

☒ Disabled ☐ Enabled

Allow recommended upper limit for shared access signature (SAS) expiry interval ⓘ

☒ Disabled ☐ Enabled

Default to Microsoft Entra authorization in the Azure portal ⓘ

☒ Disabled ☐ Enabled

Minimum TLS version ⓘ

Default to Microsoft Entra authorization in the Azure portal ⓘ

☒ Disabled ☐ Enabled

Minimum TLS version ⓘ

Version 1.2

Permitted scope for copy operations (preview) ⓘ

From any storage account

Blob access tier (default) ⓘ

☐ Cool ☒ Hot

お客様のMicrosoft Entraテナントには、以下の表に示すユーザーが含まれています。

名前  
contoso2023 用の共有アクセス署名 (SAS) トークン  
ユーザー1  
最大権限を付与したユーザー委任SAS  
ユーザー2  
利用可能な最大権限でSASサービスを提供します  
ユーザー3  
利用可能な最大権限を持つSASアカウント  
以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

| Statements                              | Yes                   | No                    |
|-----------------------------------------|-----------------------|-----------------------|
| User1 can access the content in cont1.  | <input type="radio"/> | <input type="radio"/> |
| User2 can access the content in cont1.  | <input type="radio"/> | <input type="radio"/> |
| User3 can access the content in share1. | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                              | Yes                              | No                               |
|-----------------------------------------|----------------------------------|----------------------------------|
| User1 can access the content in cont1.  | <input checked="" type="radio"/> | <input type="radio"/>            |
| User2 can access the content in cont1.  | <input type="radio"/>            | <input checked="" type="radio"/> |
| User3 can access the content in share1. | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

Yes, No, No


Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

The decisive setting is Allow storage account key access = Disabled , which disallows Shared Key authorization for the storage account. Microsoft states that when Shared Key access is disallowed, Azure Storage rejects requests authorized with account access keys, and only requests authorized with Microsoft Entra ID succeed. Microsoft's SAS behavior table is explicit: a User delegation SAS uses Microsoft Entra ID and is permitted, while a Service SAS and an Account SAS use Shared Key and are denied when AllowSharedKeyAccess is false.

Therefore, User1 can access cont1 because a user delegation SAS is Microsoft Entra-authorized and is valid for Blob Storage access. User2 cannot access cont1 because a service SAS is secured by the storage account key, which is blocked. User3 cannot access share1 because an account SAS is also secured by the storage account key and is denied for all Azure Storage services when Shared Key access is disabled. The anonymous blob access, TLS, and portal authorization settings do not override SAS authorization behavior here. Study Guide reference: AZ-104 Implement and manage storage > configure storage security, shared access signatures, access keys, Azure Blob Storage, and Azure Files.

質問: 9

次の図に示すようにピアリングが構成されています。



ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Hosts on vNET6 can communicate with hosts on [answer choice].

vNET6 only

vNET6 only

vNET6 and vNET 1 only

vNET6, vNET1, and vNET2 only

all the virtual networks in the subscription

To change the status of the peering connection to vNET1 to **Connected**, you must first [answer choice].

delete peering1

add a service endpoint

add a subnet

delete peering1

modify the address space

正解:

Answer Area

Hosts on vNET6 can communicate with hosts on [answer choice].

vNET6 only

vNET6 only

vNET6 and vNET 1 only

vNET6, vNET1, and vNET2 only

all the virtual networks in the subscription

To change the status of the peering connection to vNET1 to **Connected**, you must first [answer choice].

delete peering1

add a service endpoint

add a subnet

delete peering1

modify the address space

Explanation:

Answer Area

Hosts on vNET6 can communicate with hosts on [answer choice].

vNET6 only

To change the status of the peering connection to vNET1 to **Connected**, you must first [answer choice].

delete peering1

In Azure, Virtual Network Peering connects two virtual networks (VNETs) and allows traffic between them using private IP addresses, traversing the Microsoft backbone network rather than the public internet. For peering to function correctly, both sides of the connection must have a valid, reciprocal configuration (peering must be created from each VNet and accepted).

From the exhibit:

- \* vNET6 has two peering configurations:
- \* peering1 # vNET1 (Gateway Transit = Enabled)
- \* peering2 # vNET2 (Gateway Transit = Disabled)Both are currently showing a status of "Disconnected", meaning that the remote peering (from vNET1 and vNET2 back to vNET6) hasn't been set up or accepted yet.

Because the peering links are disconnected, vNET6 cannot communicate with any other VNet-only internal communication within vNET6 is possible.

To fix the peering connection status for vNET1, you cannot directly edit an incomplete or mismatched configuration. Microsoft's official documentation ( " Azure Virtual Network peering - requirements and constraints " ) states that you must delete and recreate the peering if the configuration is invalid or disconnected. Therefore, to establish a working connection to vNET1, you must first delete peering1 and then reconfigure it correctly on both VNETs.

質問: 10

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、指定された目標を達成する可能性のある独自のソリューションが含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

Adatumという名前のAzure Active Directory (Azure AD) テナントとSubscription1という名前のAzureサブスクリプションがあります。Adatumには、Developersという名前のグループが含まれています。Subscription1には、Devという名前のリソースグループが含まれています。

開発者グループに、DevリソースグループでAzureロジックアプリを作成する機能を提供する必要があります。

解決策：開発では、開発者グループにコントリビューターの役割を割り当てます。

これは目標を達成していますか？

A. はい

B. いいえ

正解: [\(正解を表示します\)](#)

The Contributor role grants the ability to create and manage all types of Azure resources, including logic apps. Assigning this role to the Developers group on the Dev resource group will allow them to create logic apps in that scope. Then, References: [Built-in roles for Azure resources] [Azure Logic Apps permissions and access control]

質問: 11

次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name  | Type            | Resource group |
|-------|-----------------|----------------|
| VNET1 | Virtual network | RG1            |
| VM1   | Virtual machine | RG1            |

ポリシーの適用が有効になっている 許可されていないリソースの種類」Azure ポリシーが RG1 に割り当てられており、次のパラメーターが使用されます。

Microsoft.Network/仮想ネットワーク

Microsoft.Compute/仮想マシン

RG1 で、VNET1 に接続する VM2 という新しい仮想マシンを作成する必要があります。まず何をすべきでしょうか？

- A. Azure Resource Manager テンプレートを作成します。
- B. VNET1にサブネットを追加します。
- C. ポリシーから Microsoft.Network/virtualNetworks を削除します。
- D. ポリシーから Microsoft.Compute/virtualMachines を削除します。

正解: [\(正解を表示します\)](#)

Option A (Create an Azure Resource Manager template): This wouldn ' t circumvent the policy enforcement. Even with a template, you cannot create resources that the policy explicitly denies.

Option B (Add a subnet to VNET1): Adding a subnet does not address the policy restriction on creating virtual machines. Also, the existing VNET1 can already have multiple subnets.

Option C (Remove Microsoft.Network/virtualNetworks from the policy): This isn ' t necessary because you ' re not trying to create a new virtual network; you are connecting to an existing one, VNET1.

Option D (Remove Microsoft.Compute/virtualMachines from the policy): This is the correct action because it directly addresses the restriction that is preventing you from creating a new virtual machine in RG1. Removing the virtual machine resource type from the not allowed list in the policy will enable you to create VM2. Remember, changes to policies might take a few minutes to propagate. After updating the policy, you should be able to create the new virtual machine VM2 and connect it to VNET1.

質問: 12

次の表に示すリソースが含まれる Subscription1 という名前の Azure サブスクリプションがあります。

| Name   | Type                    | Region       | Resource group |
|--------|-------------------------|--------------|----------------|
| RG1    | Resource group          | West Europe  | Not applicable |
| RG2    | Resource group          | North Europe | Not applicable |
| Vault1 | Recovery Services vault | West Europe  | RG1            |

次の表に示すように、Subscription1 に仮想マシンを作成します。

| Name | Resource group | Region       | Operating system    |
|------|----------------|--------------|---------------------|
| VM1  | RG1            | West Europe  | Windows Server 2016 |
| VM2  | RG1            | North Europe | Windows Server 2016 |
| VM3  | RG2            | West Europe  | Windows Server 2016 |
| VMA  | RG1            | West Europe  | Ubuntu Server 18.04 |
| VMB  | RG1            | North Europe | Ubuntu Server 18.04 |
| VMC  | RG2            | West Europe  | Ubuntu Server 18.04 |

できるだけ多くの仮想マシンのバックアップに Vault1 を使用する予定です。  
どの仮想マシンを Vault1 にバックアップできますか？

- A. VM1、VM3、VMA、VMCのみ
- B. VM1とVM3のみ
- C. VM1、VM2、VM3、VMA、VMB、VMC
- D. VM1のみ
- E. VM3 および VMC のみ

正解: [\(正解を表示します\)](#)

To create a vault to protect virtual machines, the vault must be in the same region as the virtual machines. If you have virtual machines in several regions, create a Recovery Services vault in each region.

References:

<https://docs.microsoft.com/bs-cyrl-ba/azure/backup/backup-create-rs-vault>

質問: 13

Windows Server 2019 を実行する VM1 という名前の Azure VM を作成します。  
VM1 は図に示すように構成されています (図タブをクリックします)。

### C. VM1 のスナップショットをキャプチャします。

D. VM1 に接続します。

正解: ([正解を表示します](#))

Scenario Overview:

You have a Windows Server 2019 virtual machine (VM1) that has been created in Azure but is currently stopped (deallocated), as indicated by the exhibit (typical of this question from the AZ-104 pool). You need to enable Desired State Configuration (DSC) on this VM.

Understanding Azure Desired State Configuration (DSC)

Azure Desired State Configuration (DSC) is a management platform used to define and maintain configuration consistency for Azure virtual machines and other systems. In Azure, DSC can be managed via Azure Automation State Configuration, or locally using PowerShell DSC on the virtual machine.

To enable DSC on an Azure VM:

The VM must be running and accessible so that the DSC extension can be installed.

The Azure VM Agent must be operational (it runs only when the VM is started).

Configuration scripts or extensions (like Microsoft.Powershell.DSC) can only be applied to an active VM.

Why "Start VM1" is the Correct First Step

From the Microsoft Azure Administrator documentation (PowerShell DSC Extension for Azure VMs):

"Before you can enable the Azure Desired State Configuration extension on a virtual machine, the VM must be in the running state. The Azure VM Agent communicates with Azure to install and configure the DSC extension, and this agent only runs when the VM is powered on." This means that if VM1 is stopped (deallocated), Azure cannot deploy or enable the DSC extension, because it requires the VM to be active and connected to the Azure platform fabric.

Why the Other Options Are Incorrect

A). Configure a DNS name for VM1:

DNS naming allows external or internal connectivity but is not a prerequisite for enabling DSC. DSC configuration relies on the Azure VM Agent, not on DNS name resolution.

C). Capture a snapshot of VM1:

Snapshots are used for backup or replication, not for enabling configuration management features such as DSC.

D). Connect to VM1:

You cannot connect to VM1 (RDP/SSH) until it is started. Starting the VM is a prerequisite before any configuration management operation.

Key Documentation Reference (Microsoft Learn):

"The PowerShell Desired State Configuration (DSC) extension for Windows is used to configure Windows VMs. The extension requires the VM to be running so the Azure VM Agent can install and apply configurations."

- Source: Microsoft Learn - Azure Virtual Machine extensions and features overview / PowerShell DSC extension

質問: 14

10 個の Azure Web アプリをデプロイするために使用される、Template1 という名前のデプロイ テンプレートがあります。

Tempi ate 1 を展開する前に、展開するものを特定する必要があります。ソリューションでは Azure のコストを最小限に抑える必要があります。何を特定する必要がありますか？

A. 10 個の App Service プラン

B. 5 つの Azure アプリケーション ゲートウェイ

C. 1 つの Azure アプリケーション ゲートウェイ

D. 1 つの Azure Traffic Manager

E. 1 つの App Service プラン

正解: ([正解を表示します](#))

質問: 15

注: この質問は、同じシナリオを提示する一連の質問の一部です。一連の質問にはそれぞれ、定められた目標を満たす可能性のある独自の解答が含まれています。質問セットによっては、複数の正解が存在する場合もあれば、正解がない場合もあります。

このセクションの質問に回答した後は、その質問に戻ることはできません。そのため、これらの質問はレビュー画面に表示されません。

contoso.com という名前の Microsoft Entra テナントがあります。

500 人の外部ユーザーの名前とメール アドレスを含む CSV ファイルがあります。

500 人の外部ユーザーごとに、contoso.com にゲスト ユーザー アカウントを作成する必要があります。

解決策: Azure ポータルの Microsoft Entra ID から、ユーザーを一括招待する操作を使用します。

これは目標を満たしていますか？

**A.** はい

**B.** いいえ

正解: ([正解を表示します](#))

This scenario involves inviting 500 external users (B2B guests) into your Microsoft Entra tenant (formerly Azure Active Directory). The goal is to create guest user accounts efficiently using a CSV file containing user details (names and email addresses).

# Understanding the Requirement

You must onboard external users (not internal directory users). These external accounts require B2B guest invitations, not normal user creation.

There are two supported approaches for bulk guest invitations:

Using the Microsoft Entra portal (Bulk invite users)

Using PowerShell with the New-AzureADMSInvitation cmdlet.

The Bulk invite users feature in the Entra portal is specifically designed for this exact use case - bulk creation of guest accounts using a CSV file.

# How the Bulk Invite Process Works

When you go to Microsoft Entra ID # Users # Bulk operations # Bulk invite, you can:

Upload a CSV file containing columns like Email address, Display name, and First/Last names.

The portal validates the file format and initiates invitations for all 500 external users.

Each external user receives an invitation email and is added as a guest (UserType = Guest) in the directory.

Once accepted, these users can be assigned roles, access resources, or be part of Azure AD groups and applications.

# Why This Meets the Goal

The solution correctly leverages the Bulk invite users functionality designed for mass external user onboarding.

It satisfies all the requirements:

Creates guest accounts (not internal users).

Uses a CSV file, allowing batch import.

Achieves the goal using Microsoft Entra ID portal, without needing PowerShell scripting.

# Reference from Microsoft Documentation (Microsoft Learn - Verified Source)

"You can use the Azure portal to invite multiple external users to your organization in one go using the Bulk invite feature. This process imports a CSV file containing user details and sends invitations to all users automatically." (Source: Microsoft Learn - Bulk invite B2B users to Microsoft Entra ID)

質問: **16**

storage1 という名前の Azure ストレージ アカウントがあります。

User1 という名前のユーザーが storage1 のストレージ アカウント キーを一覧表示および再生成できるようにする必要があります。

解決策: ストレージ アカウント キー オペレーターのサービス ロールを User1 に割り当てます。

これは目標を満たしていますか？

- A. はい
- B. いいえ

正解: (正解を表示します)

In Azure Role-Based Access Control (RBAC), the Storage Account Key Operator Service Role allows a user to list and regenerate access keys for a storage account but does not grant permission to access the data itself.

The Storage Account Key Operator Service Role has the following key permissions:

- \* Microsoft.Storage/storageAccounts/listkeys/action
- \* Microsoft.Storage/storageAccounts/regeneratekey/action
- \* Microsoft.Storage/storageAccounts/read

This role is typically used by automation or service accounts that need to manage storage account keys without having full access to the data stored within the account (blobs, files, queues, etc.).

In the scenario, User1 only needs to list and regenerate storage account keys for storage1, not perform data operations or modify configurations. The Storage Account Key Operator Service Role provides precisely these permissions, making it the correct and minimal-privilege choice for this task.

Thus, assigning this role to User1 meets the goal as per Azure RBAC guidelines and Microsoft's least privilege principle documented in the Azure Administrator exam guide.

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 17

storage1 という名前の Azure ストレージ アカウントと、次の表に示すユーザーを含む Azure サブスクリプションがあります。

| Name  | Member of |
|-------|-----------|
| User1 | Group1    |
| User2 | Group2    |
| User3 | Group1    |

ストレージ 1 を監視し、次の表に示す信号に関する電子メール通知を構成する予定です。

| Name                   | Type         | Users to notify         |
|------------------------|--------------|-------------------------|
| Ingress                | Metric       | User1 and User3 only    |
| Egress                 | Metric       | User1 only              |
| Delete storage account | Activity log | User1, User2, and User3 |
| Restore blob ranges    | Activity log | User1 and User3 only    |

計画された監視に必要なアラート ルールとアクション グループの最小数を特定する必要があります。

アラート ルールとアクション グループをいくつ特定する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Alert rules:

|   |   |
|---|---|
|   | ▼ |
| 1 |   |
| 2 |   |
| 3 |   |
| 4 |   |

Action groups:

|                                                                                   |           |   |
|-----------------------------------------------------------------------------------|-----------|---|
|  | Microsoft | ▼ |
| 1                                                                                 |           |   |
| 2                                                                                 |           |   |
| 3                                                                                 |           |   |
| 4                                                                                 |           |   |

正解:

Alert rules:

|   |   |
|---|---|
|   | ▼ |
| 1 |   |
| 2 |   |
| 3 |   |
| 4 |   |

Action groups:

|   |   |
|---|---|
|   | ▼ |
| 1 |   |
| 2 |   |
| 3 |   |
| 4 |   |

 Microsoft

Explanation:

Box 1 : 4

As there are 4 distinct set of resource types (Ingress, Egress, Delete storage account, Restore blob ranges), so you need 4 alert rules. In one alert rule you can ' t specify different type of resources to monitor. So you need 4 alert rules.

Box 2 : 3

There are 3 distinct set of " Users to notify " as (User 1 and User 3), (User1 only), and (User1, User2, and User3). You can ' t set the action group based on existing group (Group1 and Group2) as there is no specific group for User1 only. So you need to create 3 action group.

Answer Area

Alert rules:

|   |   |
|---|---|
|   | ▼ |
| 1 |   |
| 2 |   |
| 3 |   |
| 4 |   |

Action groups:

|   |   |
|---|---|
|   | ▼ |
| 1 |   |
| 2 |   |
| 3 |   |
| 4 |   |

 Microsoft

Reference:  
<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/action-groups>

質問: 18

Azure Active Directory (Azure AD) テナントがあります。  
Azure Active Directory 管理センターで一括削除を使用して、複数のユーザーを削除することを計画しています。  
一括削除用のファイルを作成してアップロードする必要があります。  
どのユーザー属性をファイルに含める必要がありますか？

- A. 各ユーザーのユーザープリンシパル名と使用場所のみ
- B. 各ユーザーのユーザープリンシパル名のみ
- C. 各ユーザーの表示名のみ
- D. 各ユーザーの表示名と使用場所のみ
- E. 各ユーザーの表示名とユーザープリンシパル名のみ

正解: (正解を表示します)

To perform a bulk delete of users in Azure Active Directory, you need to create and upload a CSV file that contains the list of users to be deleted. The file should include the user principal name (UPN) of each user only. Therefore, the answer is B. The user principal name of each user only. When you use the bulk delete feature in the Azure Active Directory admin center, you need to specify the UPN for each user that you want to delete. The UPN is a unique identifier for each user in Azure AD and is the primary way that Azure AD identifies and manages user accounts. Including additional attributes like the display name or usage location is not required for the bulk delete operation, as the UPN is the only mandatory attribute for the user account. However, you may include additional attributes in the CSV file if you want to keep track of the metadata associated with each user account.

質問: 19

次の表に示すパブリック IP アドレスを使用する Azure サブスクリプションがあります。

| Name | IP version | SKU      | IP address assignment | Availability zone |
|------|------------|----------|-----------------------|-------------------|
| IP1  | IPv6       | Basic    | Static                | Not applicable    |
| IP2  | IPv6       | Basic    | Dynamic               | Not applicable    |
| IP3  | IPv6       | Standard | Static                | Zone-redundant    |

パブリック Azure Standard Load Balancer を作成する必要があります。  
どのパブリック IP アドレスを使用できますか？

- A. IP1とIP3のみ
- B. IP1、IP2、IP3
- C. IP2のみ
- D. IP3のみ

正解: [\(正解を表示します\)](#)

A Basic Load Balancer can use the Basic SKU Public IP address 's, but a Standard load balancer requires a Standard SKU Public IP address.

Excerpt from link below:

The standard SKU is required if you associate the address to a standard load balancer. For more information about standard load balancers, see Azure load balancer standard SKU.

<https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/virtual-network-public-ip-address> Excerpt from link below:

Key scenarios that you can accomplish using Azure Standard Load Balancer include:

-Enable support for load-balancing of IPv6.

<https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-overview#why-use-azure-load-balancer>

質問: 20

Azureサブスクリプションに、storage1というストレージアカウントが含まれています。storage1アカウントには、containerというコンテナが含まれています。container1へのアクセスを構成する必要があります。ソリューションは以下の要件を満たす必要があります。

\* 読み取りアクセスのみ許可

\* HTTP プロトコルと HTTPS プロトコルの両方を許可します。

\* コンテナ内のすべてのコンテンツにアクセス権限を適用する

何を使うべきでしょうか？

- A. アクセスポリシー
- B. 共有アクセス署名 (SAS)
- C. Azure コンテンツ配信ネットワーク (CDN)
- D. アクセスキー

正解: B [\(コメントを發表する\)](#)

According to the Microsoft documentation, a shared access signature (SAS) is a URI that grants restricted access rights to Azure Storage resources. You can provide a SAS to clients who don't otherwise have access to your storage account, and delegate access to them for a specified time period and with a specified set of permissions.

A SAS can be used to grant read-only access to a container and its blobs, as well as specify the allowed protocols (HTTP or HTTPS) and the start and expiry time of the access. For more information about creating and using SAS, see Using shared access signatures (SAS).

An access policy is not the correct answer because it is used to define a set of permissions and a time period for a container or a queue, but it does not grant access by itself. An access policy must be associated with a SAS to take effect. For more information about access policies, see Manage stored access policies for containers and queues.

Azure Content Delivery Network (CDN) is not the correct answer because it is used to cache and deliver content from Azure Storage or other sources, but it does not control the access permissions to the content. For more information about Azure CDN, see [What is Azure Content Delivery Network?].

Access keys are not the correct answer because they are used to authenticate requests to Azure Storage from any client, but they do not limit the access permissions or the protocols. Using access keys also exposes your storage account to potential unauthorized access if the keys are compromised. For more information about access keys, see [Manage storage account access keys].

質問: 21

Microsoft 365 テナントと、contoso.com という名前の Azure Active Directory (Azure AD) テナントがあります。  
User1、User2、User3 という 3 人のユーザーに、Library1 という一時的な Microsoft SharePoint ドキュメント ライブラリへのアクセスを許可する予定です。  
ユーザー用のグループを作成する必要があります。ソリューションでは、グループが180日後に自動的に削除されるようにする必要があります。  
どの 2 つのグループを作成する必要がありますか？それぞれの正解は完全なソリューションを示します。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

- A. 割り当てられたメンバーシップタイプを使用するセキュリティグループ
- B. 割り当てられたメンバーシップ タイプを使用する Office 365 グループ
- C. 動的ユーザー メンバーシップ タイプを使用する Office 365 グループ
- D. 動的ユーザー メンバーシップ タイプを使用するセキュリティ グループ
- E. 動的デバイスメンバーシップタイプを使用するセキュリティグループ

正解: (正解を表示します)

You can set expiration policy only for Office 365 groups in Azure Active Directory (Azure AD).  
Note: With the increase in usage of Office 365 Groups, administrators and users need a way to clean up unused groups. Expiration policies can help remove inactive groups from the system and make things cleaner.  
When a group expires, all of its associated services (the mailbox, Planner, SharePoint site, etc.) are also deleted.  
You can set up a rule for dynamic membership on security groups or Office 365 groups.

質問: 22

次の表に示す仮想ネットワークを含む Azure サブスクリプションがあります。

| Name  | Azure region | Resource group |
|-------|--------------|----------------|
| VNET1 | West US      | RG1            |
| VNET2 | Central US   | RG1            |
| VNET3 | Central US   | RG2            |
| VNET4 | West US      | RG2            |

- 米国西部 Azure リージョンの RG1 に AF1 という名前の Azure ファイアウォールをデプロイする必要があります。  
AF1 をどの仮想ネットワークにデプロイできますか？
- A. VNET1のみ
  - B. VNET1 と VNET2 のみ
  - C. VNET1 と VNET4 のみ
  - D. VNET1、VNET2、および VNET4 のみ
  - E. VNET1、VNET2、VNET3、および VNET4


正解: (正解を表示します)

Azure Firewall must be deployed in the same Azure region as the virtual network and into a dedicated subnet named AzureFirewallSubnet.  
From the table:  
\* VNET1 # West US # RG1 #  
\* VNET4 # West US # RG2 #  
\* VNET2 # Central US #  
\* VNET3 # Central US #  
Resource group location does not restrict deployment-only region alignment matters.

Microsoft documentation confirms:  
"Azure Firewall must be deployed into a virtual network in the same region."


質問: 23

次の図に示すストレージ アカウントを含む Azure サブスクリプションがあります。



ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

正解:



Explanation:



質問: 24

次の表に示すAzure Webアプリを作成する予定です。

| Name    | Runtime stack |
|---------|---------------|
| WebApp1 | .NET 6 (LTS)  |
| WebApp2 | ASP.NET V4.8  |
| WebApp3 | PHP 8.1       |
| WebApp4 | Python 3.11   |

Webアプリ用に作成する必要があるApp Serviceプランの最小数はいくつですか？

- A. 1
- B. 2
- C. 3
- D. 4

正解: ([正解を表示します](#))

NET Core 3.0: Windows and Linux ASP .NET V4.7: Windows only PHP 7.3: Windows and Linux Ruby 2.6: Linux only Also, you can't use Windows and Linux Apps in the same App Service Plan, because when you create a new App Service plan you have to choose the OS type. You can ' t mix Windows and Linux apps in the same App Service plan. So, you need 2 ASPs. Reference: <https://docs.microsoft.com/en-us/azure/app- service/overview>

質問: 25

技術要件とユーザー要件を満たすようにデバイス設定を構成する必要があります。

どの 2 つの設定を変更する必要がありますか? 回答するには、回答領域で適切な設定を選択してください。

Answer Area

 Save

 Discard

Users may join devices to Azure AD ⓘ

All

Selected

None

Selected

No member selected

Additional local administrators on Azure AD joined devices ⓘ

Selected

None

Selected

No member selected

Users may register their devices with Azure AD ⓘ

All

None

Require Multi-Factor Auth to join devices ⓘ

Yes

No

Maximum number of devices per user ⓘ

50

Users may sync settings and app data across devices ⓘ

All

Selected

None

Selected

No member selected

正解:

Users may join devices to Azure AD ⓘ

All

Selected

None

Selected

No member selected

Additional local administrators on Azure AD joined devices ⓘ

Selected

None

Selected

No member selected

Users may register their devices with Azure AD ⓘ

All

None

Require Multi-Factor Auth to join devices ⓘ

Yes

No

Maximum number of devices per user ⓘ

50

Users may sync settings and app data across devices ⓘ

All

Selected

None

Selected

No member selected

Explanation:

Save

Discard

Users may join devices to Azure AD ⓘ

AllSelectedNone

Selected

No member selected

Additional local administrators on Azure AD joined devices ⓘ

SelectedNone

Selected

No member selected

Users may register their devices with Azure AD ⓘ

AllNone

Require Multi-Factor Auth to join devices ⓘ

YesNo

Maximum number of devices per user ⓘ

50

Users may sync settings and app data across devices ⓘ

AllSelectedNone

Box 1: Selected

Only selected users should be able to join devices

Box 2: Yes

Require Multi-Factor Auth to join devices.

From scenario:

Ensure that only users who are part of a group named Pilot can join devices to Azure AD Ensure that when users join devices to Azure Active Directory (Azure AD), the users use a mobile phone to verify their identity.

質問: 26

VM1 という名前の Azure 仮想マシンがあります。

Azure Backup を使用して、Backup1 という名前の VM1 のバックアップを作成します。

Backup1 を作成した後、VM1 に対して次の変更を実行します。

VM1 のサイズを変更します。

Budget.xls という名前のファイルを Data という名前のフォルダーにコピーします。

組み込みの管理者アカウントのパスワードをリセットします。

VM1 にデータ ディスクを追加します。

管理者は、既存の置き換えオプションを使用して、Backup1 から VM1 を復元します。

VM1 へのすべての変更が復元されていることを確認する必要があります。

どの変更を再度実行する必要がありますか？

- A. VM1 のサイズを変更します。
- B. データ ディスクを追加します。
- C. 組み込みの管理者アカウントのパスワードをリセットします。
- D. Budget.xls をデータにコピーします。

正解: B (コメントを发表する)

When you perform a VM restore using Azure Backup with the "Replace existing" option:

- \* Azure Backup replaces the existing VM's operating system disk and configuration with what was captured at the time of the backup.
- \* Data disks added after the backup are not part of the recovery point and are therefore not restored.

Other post-backup changes:

- \* VM size change: Automatically reverted to the size captured in the backup configuration.
- \* File addition (Budget.xls): Not present, because only data up to the backup time is restored.
- \* Password reset: Resetting password does not persist because OS disk from backup overwrites current one.However, when restoring, you can always reset credentials later - but that is expected behavior.

The only configuration that requires manual re-creation after restore is any new data disk added after the backup snapshot.

質問: 27

次の表に示すパブリック IP アドレスを持つ Azure サブスクリプションがあります。

| Name | IP version | SKU      | Tier     | IP address assignment |
|------|------------|----------|----------|-----------------------|
| IP1  | IPv4       | Standard | Regional | Static                |
| IP2  | IPv4       | Standard | Global   | Static                |
| IP3  | IPv4       | Basic    | Regional | Dynamic               |
| IP4  | IPv4       | Basic    | Regional | Static                |
| IP5  | IPv6       | Basic    | Regional | Dynamic               |

Bastion1 という名前の Azure Bastion Basic SKU ホストをデプロイする予定です。

Bastion1 に使用できる IP アドレスはどれですか？

- A. IP1のみ
- B. IP1とIP2のみ
- C. IP3、IP4、およびIPSのみ
- D. IP1、IP2、IP4、IP5のみ
- E. IP1、IP2、IP3、IP4、およびIPS

正解: A (コメントを发表する)

When deploying an Azure Bastion host, the network configuration must meet specific requirements depending on the Bastion SKU. For a Bastion Basic SKU, the associated public IP address must fulfill all of the following Azure-defined conditions:

SKU requirement: The Bastion Basic host requires a Standard public IP address.

Tier restriction: The IP must be Regional, not Global, because Bastion operates within a specific Azure region and cannot use global IPs.

Assignment requirement: The IP must use Static allocation, as Bastion requires a reserved, unchanging public IP for consistent access and DNS resolution.

Supported IP version: Azure Bastion supports IPv4 addresses only; IPv6 addresses are not supported.

From the table provided:

IP1 - Standard, Regional, Static, IPv4 # # Meets all requirements.

IP2 - Standard, Global, Static # # Global tier is not supported.

IP3 - Basic, Dynamic # # Bastion requires Standard SKU and Static IP.

IP4 - Basic, Static # # SKU must be Standard.

IP5 - IPv6, Dynamic # # IPv6 and dynamic IPs are unsupported.

Therefore, only IP1 meets the Microsoft Azure Bastion Basic deployment prerequisites.

Reference (Azure Administrator Docs Extract):

Azure Bastion hosts require a Standard, Regional, Static, IPv4 public IP address. Basic or Global IPs are not supported for deployment.

質問: 28

次の図に示すように、Azure ストレージ アカウントがあります。

| Storage accounts                                                                            |                 |                 |             |            |           |                |             |                    |
|---------------------------------------------------------------------------------------------|-----------------|-----------------|-------------|------------|-----------|----------------|-------------|--------------------|
| Contoso                                                                                     |                 |                 |             |            |           |                |             |                    |
| + Add Edit columns Refresh Assign Tags Delete                                               |                 |                 |             |            |           |                |             |                    |
| Subscriptions: All 2 selected - Don't see a subscription? Switch directories                |                 |                 |             |            |           |                |             |                    |
| Filter by name... All subscriptions All resource groups All types All locations No grouping |                 |                 |             |            |           |                |             |                    |
| 3 items                                                                                     |                 |                 |             |            |           |                |             |                    |
| <input type="checkbox"/>                                                                    | NAME            | TYPE            | KIND        | RESOURCE   | LOCATION  | SUBSCRIPTI...  | ACCESS T... | REPLICAT....       |
| <input type="checkbox"/>                                                                    | storageaccount1 | Storage account | Storage     | ContosoRG1 | EastUS    | Subscription 1 | -           | Read-access ge...  |
| <input type="checkbox"/>                                                                    | storageaccount2 | Storage account | StorageV2   | ContosoRG1 | CentralUS | Subscription 1 | Host        | Geo-redundant...   |
| <input type="checkbox"/>                                                                    | storageaccount3 | Storage account | BlobStorage | ContosoRG1 | EastUS    | Subscription 1 | Host        | Locally-redund.... |

ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

You can use [answer choice] for Azure Table Storage.

storageaccount1 only

storageaccount2 only

storageaccount3 only

storageaccount1 and storageaccount2 only

storageaccount2 and storageaccount3 only

You can use [answer choice] for Azure Blob storage.

storageaccount3 only

storageaccount2 and storageaccount3 only

storageaccount1 and storageaccount3 only

all the storage accounts

正解:

Answer Area

You can use [answer choice] for Azure Table Storage.

storageaccount1 only  
storageaccount2 only  
storageaccount3 only  
storageaccount1 and storageaccount2 only  
storageaccount2 and storageaccount3 only

You can use [answer choice] for Azure Blob storage.

storageaccount3 only  
storageaccount2 and storageaccount3 only  
storageaccount1 and storageaccount3 only  
all the storage accounts

Explanation:

You can use [answer choice] for Azure Table Storage.

storageaccount1 only  
storageaccount2 only  
storageaccount3 only  
storageaccount1 and storageaccount2 only  
storageaccount2 and storageaccount3 only

You can use [answer choice] for Azure Blob storage.

storageaccount3 only  
storageaccount2 and storageaccount3 only  
storageaccount1 and storageaccount3 only  
all the storage accounts

Box 1: storageaccount1 and storageaccount2 only

Box 2: All the storage accounts

Note: The three different storage account options are: General-purpose v2 (GPv2) accounts, General-purpose v1 (GPv1) accounts, and Blob storage accounts.

General-purpose v2 (GPv2) accounts are storage accounts that support all of the latest features for blobs, files, queues, and tables.

Blob storage accounts support all the same block blob features as GPv2, but are limited to supporting only block blobs.

General-purpose v1 (GPv1) accounts provide access to all Azure Storage services, but may not have the latest features or the lowest per gigabyte pricing.

References: <https://docs.microsoft.com/en-us/azure/storage/common/storage-account-options>

質問: 29

Azure ストレージ アカウントを作成します。

ストレージ アカウントに 10 個の BLOB コンテナを追加する予定です。

コンテナの 1 つでは、保存されているデータを暗号化するために別のキーを使用する必要があります。

コンテナを作成する前に何をすべきでしょうか？

- A. 最小 TLS バージョンを変更します。
- B. 暗号化スコープを作成します。
- C. 共有アクセス署名 (SAS) を生成します。
- D. アクセス キーをローテーションします。

正解: [\(正解を表示します\)](#)

Azure Storage encrypts all data at rest using Storage Service Encryption (SSE) by default. However, if you need to use a different encryption key for specific containers or blobs within the same storage account, you must create an encryption scope.

An encryption scope defines a boundary within the storage account where data encryption is handled by a unique customer-managed key (CMK) or Microsoft-managed key. You can then associate this encryption scope with a specific container or even individual blobs, allowing flexible key management across data sets.

The Microsoft Azure Storage documentation explains that encryption scopes allow you to:

- Use different encryption keys for different containers or blobs.
- Rotate or manage keys independently for compliance or separation-of-duty requirements.
- Support encryption at container creation by assigning a specific scope.

Other options are incorrect:

- Modifying the TLS version affects network security, not encryption keys.
- Generating a SAS defines access tokens, not encryption behavior.
- Rotating access keys re-generates account keys for security, but it does not create a new encryption key for specific containers.

Hence, before creating the container, you must first create an encryption scope and then assign it to that container to ensure it uses a different encryption key.

質問: 30

お客様はMicrosoft Entraテナントをお持ちです。  
新しい Microsoft 365 グループが作成される際に、グループ名が以下のように自動的にフォーマットされるようにする必要があります。  
<部門グループ名>。

Microsoft Entra 管理センターで、どの 3 つの操作を順番に実行する必要がありますか？回答するには、操作の一覧から適切な操作を回答欄に移動し、正しい順序に並べ替えてください。

Actions

Customize the company branding.

Set Add suffix to String.

Set Add suffix to Attribute.

Set Add prefix to String.

Create a group naming policy.

Set Add prefix to Attribute.

Set Select type to Department.

Answer Area

正解:

Actions

Customize the company branding.

Set Add suffix to String.

Set Add suffix to Attribute.

Set Add prefix to String.

Create a group naming policy.

Set Add prefix to Attribute.

Set Select type to Department.

Answer Area

Create a group naming policy.

Set Add prefix to Attribute.

Set Select type to Department.

Explanation:

Actions

Customize the company branding.

Set Add suffix to String.

Set Add suffix to Attribute.

Set Add prefix to String.

Answer Area

1 Create a group naming policy.

2 Set Add prefix to Attribute.

3 Set Select type to Department.

Microsoft Entra ID (formerly Azure Active Directory) supports group naming policies to automatically enforce consistent and compliant naming conventions for Microsoft 365 groups. The purpose of this policy is to standardize names, prevent conflicts, and include metadata such as department, location, or purpose.

According to the Microsoft Entra Administrator documentation and the AZ-104 official study guide ("Manage Azure Identities and Governance"), the steps to implement such a policy are:

Step 1: Create a Group Naming Policy

Navigate to:

Microsoft Entra admin center # Groups # Naming policy.

Here, you can define global naming conventions that apply whenever new Microsoft 365 groups or security groups are created.

Creating the policy allows you to specify prefixes and suffixes that automatically appear in group names based on attributes or fixed strings.

Step 2: Set Add Prefix to Attribute

Once the naming policy is created, select the Add prefix option.

You can choose between two prefix types:

String: A fixed word or phrase (for example, "Corp\_")

Attribute: A user or group attribute dynamically pulled from Microsoft Entra ID (for example, Department or CompanyName).

To achieve the naming format < Department > < Group name > , you must choose Attribute as the prefix type, since it will automatically insert the department value of the creator or group owner.


Step 3: Set Select Type to Department

After choosing Attribute as the prefix, you select which attribute will be used.

Choose Department to ensure the prefix dynamically reflects the department name of the group creator.

This results in an automatically generated group name like:  
FinanceSales, HRRecruiting, etc.  
If desired, you could later add a suffix (like "\_Group"), but for this specific scenario, only the prefix (Department) is required.  
Official Microsoft Documentation Extract (Summarized):  
"A naming policy can consist of prefixes or suffixes that include fixed strings or user attributes, such as  
[Department], [Company] , or [Office].  
To configure, create a group naming policy and add prefix or suffix elements based on attributes." (Reference: Microsoft Learn - Configure naming policy for Microsoft 365 groups in Azure Active Directory)

質問: 31  
次の図に示すように、Microsoft Entra テナントが構成されています。



テナントには、次の表に示す ID が含まれます。

| Name   | Type                |
|--------|---------------------|
| User1  | User account        |
| Group1 | Security group      |
| Group2 | Microsoft 365 group |

Microsoft Fabric ライセンスを購入します。  
どの ID にライセンスを割り当てることができますか？  
A. ユーザー1とグループ1のみ  
B. ユーザー1、グループ1、グループ2  
C. ユーザー1のみ

D. ユーザー1とグループ2のみ

正解: C ([コメントを发表する](#))

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%**ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 32

plan1 という App Service プランと、webapp1 という Azure Web アプリを作成しました。すると、ステージング スロットを作成するオプションが利用できないことがわかりました。plan1 用のステージング スロットを作成する必要があります。

まず何をすべきでしょうか？

A. webapp1 から、アプリケーション設定を変更します。

B. webapp1 からカスタム ドメインを追加します。

C. plan1 から、App Service プランをスケールアップします。

D. plan1 から、App Service プランをスケールアウトします。

正解: ([正解を表示します](#))

The app must be running in the Standard, Premium, or Isolated tier in order for you to enable multiple deployment slots. If the app isn't already in the Standard, Premium, or Isolated tier, you receive a message that indicates the supported tiers for enabling staged publishing. At this point, you have the option to select Upgrade and go to the Scale tab of your app before continuing.

Scale up: Get more CPU, memory, disk space, and extra features like dedicated virtual machines (VMs), custom domains and certificates, staging slots, autoscaling, and more.

Scale out: Increase the number of VM instances that run your app. You can scale out to as many as 30 instances Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots>

<https://docs.microsoft.com/en-us/azure/app-service/manage-scale-up>

質問: 33

Subscription1 という名前の Azure サブスクリプションがあります。

Subscription1 で、Alert1 という名前のアラート ルールを作成します。

Alert1 アクション グループは、次の図に示すように構成されています。

```
PS Azure:\> Get-AzureRmActionGroup

ResourceGroupName: default-activitylogalerts
GroupShortName    : AG1
Enabled           : True
EmailReceivers    : {Action1_EmailAction-}
SmsReceivers      : {Action1_SMSAction-}
WebhookReceivers  : {}
Id                : /subscriptions/a4fde29b-d56a-4f6c-8298-6c53cd0b720c/
resourceGroups/default-activitylogalerts/providers/microsoft.insights/actionGroups/ActionGroup1
Name              : ActionGroup1
Type              : Microsoft.Insights/ActionGroups
Location          : Global
Tags              : {}
```

Alert1 アラート基準は 1 分ごとにトリガーされます。

ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

注意: 正しい選択ごとに 1 ポイントが加算されます。

The number of email messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |

The number of SMS messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |

正解:

The number of email messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
|    |
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |

The number of SMS messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
|    |
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |

Explanation:

The number of email messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
|    |
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |

The number of SMS messages that Alert1 will send in an hour is [answer choice].

|    |
|----|
|    |
| 0  |
| 4  |
| 6  |
| 12 |
| 60 |



Box 1: 60

One alert per minute will trigger one email per minute.

Box 2: 12

No more than 1 SMS every 5 minutes can be send, which equals 12 per hour.

Note: Rate limiting is a suspension of notifications that occurs when too many are sent to a particular phone number, email address or device. Rate limiting ensures that alerts are manageable and actionable.

The rate limit thresholds are:

SMS: No more than 1 SMS every 5 minutes.

Voice: No more than 1 Voice call every 5 minutes.

Email: No more than 100 emails in an hour.

Other actions are not rate limited.

References:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring-overview-alerts.md>

質問: 34

次の表に示すパブリック ロード バランサーを含む Azure サブスクリプションがあります。

| Name | SKU      |
|------|----------|
| LB1  | Basic    |
| LB2  | Standard |

6台の仮想マシンを作成し、それらの仮想マシンへのリクエストを負荷分散する予定です。各ロードバランサーは3台の仮想マシンを負荷分散します。

計画されたソリューション用の仮想マシンを作成する必要があります。

仮想マシンはどのように作成すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

The virtual machines that will be load balanced by using LB1 must:

be created in the same availability set or virtual machine scale set.

be connected to the same virtual network.

be created in the same resource group.

be created in the same availability set or virtual machine scale set.

run the same operating system.

The virtual machines that will be load balanced by using LB2 must:

be connected to the same virtual network.

be connected to the same virtual network.

be created in the same resource group.

be created in the same availability set or virtual machine scale set.

run the same operating system.



Microsoft

正解:

**Answer Area**

The virtual machines that will be load balanced by using LB1 must:

- be created in the same availability set or virtual machine scale set.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

The virtual machines that will be load balanced by using LB2 must:

- be connected to the same virtual network.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

Explanation:

**Answer Area**

The virtual machines that will be load balanced by using LB1 must:

The virtual machines that will be load balanced by using LB2 must:

Azure Load Balancers are offered in two SKUs - Basic and Standard - each with distinct capabilities, scalability, and configuration requirements as documented in Microsoft Azure Administrator documentation.

\* LB1 - Basic Load Balancer (Basic SKU)The Basic Load Balancer is designed for small-scale, non- production workloads. It has certain limitations compared to the Standard SKU:

- \* The backend pool of a Basic Load Balancer can only include virtual machines in a single availability set or a single virtual machine scale set.
- \* This restriction ensures that the Basic Load Balancer maintains consistency and proper failover across VMs sharing the same availability set or scale set.
- \* It cannot span multiple availability sets or virtual networks.

Therefore, to load balance VMs using a Basic Load Balancer, those VMs must be created in the same availability set or virtual machine scale set.

# Correct Option for LB1: be created in the same availability set or virtual machine scale set.

\* LB2 - Standard Load Balancer (Standard SKU)The Standard Load Balancer is designed for production- grade workloads and supports advanced features such as zone redundancy, secure by default configuration, and high scalability.

- \* It supports backend pools composed of virtual machines from different availability zones (in the same region) and multiple virtual machine scale sets.
- \* The only requirement is that all backend virtual machines must be connected to the same virtual network.
- \* Standard Load Balancer also supports both public and internal load balancing and provides richer metrics and diagnostic capabilities.

Therefore, for LB2 (Standard), the virtual machines only need to reside within the same virtual network, regardless of their availability sets or zones.

# Correct Option for LB2: be connected to the same virtual network.

Microsoft Azure Official Extract (summarized):

"For Basic Load Balancer, backend pool VMs must belong to the same availability set or virtual machine scale set. For Standard Load Balancer, backend pool members can be across availability zones and scale sets, but must be connected to the same virtual network." (From Microsoft Azure Administrator Study Guide - Load Balancer Configuration and Comparison; Azure Docs: [azure.microsoft.com](https://azure.microsoft.com/en-us/load-balancer-skus-comparison/) > Load Balancer SKUs Comparison.)

Final Verified Answer:

# LB1: be created in the same availability set or virtual machine scale set

# LB2: be connected to the same virtual network

質問: 35

VM1 と VM2 という名前の 2 つの仮想マシンを含む Azure サブスクリプションがあり、Azure ロード バランサーを作成します。

VM1 と VM2 の間で HTTPS トラフィックを負荷分散する負荷分散ルールを作成する予定です。

負荷分散ルールを作成する前に、どの 2 つの追加の負荷分散リソースを作成する必要がありますか？

それぞれの正解は解決策の一部を示す

MOTL 正解 5 つにつき 1 ポイント獲得。

**A.** フロントエンドIPアドレス

**B.** バックエンドプール

**C.** ヘルスプローブ

**D.** 受信NATルール

**E.** 仮想ネットワーク

正解: **A,C** ([コメントを発表する](#))

To create a load balancing rule that will load balance HTTPS traffic between VM1 and VM2, you need to create two additional load balance resources: a frontend IP address and a health probe.

A frontend IP address is the IP address that the clients use to access the load balancer. It can be either public or private, depending on the type of load balancer. A frontend IP address is required for any load balancing rule1.

A health probe is used to monitor the health and availability of the backend instances. It can be either TCP, HTTP, or HTTPS, depending on the protocol of the load balancing rule. A health probe is required for any load balancing rule1.

A backend pool is a group of backend instances that receive the traffic from the load balancer. You already have a backend pool that contains VM1 and VM2, so you don't need to create another one.

An inbound NAT rule is used to forward traffic from a specific port on the frontend IP address to a specific port on a backend instance. It's not required for a load balancing rule, but it can be used to access individual instances for troubleshooting or maintenance purposes1.

A virtual network is a logical isolation of Azure resources within a region. It's not a load balance resource, but it's required for creating an internal load balancer or connecting virtual machines to a load balancer2.

質問: 36

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

AzureサブスクリプションのTrafficAnalyticsを有効にするには、Admin1という名前のAzure Active Directory (Azure AD) ユーザーに必要な役割が割り当てられていることを確認する必要があります。

解決策 :サブスクリプションレベルで所有者の役割をAdmin1に割り当てます。

これは目標を達成していますか？

**A.** はい

**B.** いいえ

正解: ([正解を表示します](#))

To enable Traffic Analytics for an Azure subscription, the user must have sufficient privileges to configure Network Watcher, NSG flow logs, and the associated Log Analytics workspace.

As per Microsoft Azure documentation, the following built-in roles can enable Traffic Analytics:

\* Owner

\* Contributor

\* Reader

\* Network Contributor

The Owner role provides full access to all resources, including the right to delegate permissions and modify configurations. Since the Owner role includes complete management capabilities for all Azure resources at the subscription level, this role absolutely meets the requirements for enabling Traffic Analytics.

The Azure Network Watcher documentation clearly states:

"To enable Traffic Analytics, your account must have any one of the following roles at the subscription scope: Owner, Contributor, Reader, or Network Contributor." Therefore, assigning the Owner role to Admin1 at the subscription level ensures Admin1 has the required permissions to enable Traffic Analytics.

質問: 37

Admin1 のユーザー要件を満たす必要があります。  
何をすべきでしょうか？

- A. [サブスクリプション] ブレードからサブスクリプションを選択し、[プロパティ] を変更します。
- B. [サブスクリプション] ブレードからサブスクリプションを選択し、[アクセス制御 (IAM)] 設定を変更します。
- C. Azure Active Directory ブレードから、プロパティを変更します。
- D. Azure Active Directory ブレードから、グループを変更します。

正解: A (コメントを發表する)

In this scenario, Contoso Ltd. must designate a new user named Admin1 as the service administrator of the Azure subscription and ensure that Admin1 receives email alerts about service outages.

In Azure, there are three classic administrative roles associated with subscriptions:

- \* Account Administrator - The individual who created the subscription and manages billing.
- \* Service Administrator - The primary contact responsible for managing services and resources in the subscription.
- \* Co-Administrator - Has the same management privileges as the Service Administrator but cannot change subscription associations.

According to Microsoft Azure Administrator documentation, to change the Service Administrator, you must perform the following steps:

- \* In the Azure portal, navigate to Subscriptions.
- \* Select the specific subscription.
- \* Under Settings, choose Properties.
- \* In the Service Administrator field, update the name and email address to that of the new administrator (in this case, Admin1).

This modification ensures that Admin1 becomes the Service Administrator, and Azure will automatically send service-related notifications and outage alerts to that user's registered email address, as defined by Microsoft' s subscription notification process.

Option B (IAM settings) is used for Role-Based Access Control (RBAC) and assigning Azure Resource Manager roles such as Owner, Contributor, or Reader. However, RBAC roles do not change the Service Administrator at the subscription level - that's only done through the Properties blade.

Options C (AAD Properties) and D (AAD Groups) are unrelated to subscription-level administrative settings.

Therefore, the verified and Microsoft-documented answer is:

# A. From the Subscriptions blade, select the subscription, and then modify the Properties.

Final Verified Answer: # A. From the Subscriptions blade, select the subscription, and then modify the Properties.

質問: 38

次の表に示す IP アドレスを使用する Cluster1 という名前の Azure Kubernetes Service (AKS) クラスターをデプロイします。

| IP address   | Assigned to             |
|--------------|-------------------------|
| 131.107.2.1  | Load balancer front end |
| 192.168.10.2 | Kubernetes DNS service  |
| 172.17.7.1   | Docker bridge address   |
| 10.0.10.11   | Kubernetes cluster node |

インターネット ユーザーに Cluster1 で実行されるアプリケーションへのアクセスを提供する必要があります。  
Ousted の DNS レコードに含めるべき IP アドレスはどれですか？

- A. 172.17.7.1

- B.** 131.107.2.1
- C.** 192.168.10.2
- D.** 10.0.10.11

正解: ([正解を表示します](#))

When any internet user will try to access the cluster which is behind a load balancer, traffic will first hit to load balancer front end IP. So in the DNS configuration you have to provide the IP address of the load balancer.

Reference:

<https://stackoverflow.com/questions/43660490/giving-a-dns-name-to-azure-load-balancer>

質問: **39**

VNet1 という仮想ネットワークを含む、Subscription1 という Azure サブスクリプションがあります。VNet1 は、RG1 というリソース グループに属しています。

Subscription1 には User1 という名前のユーザーがいます。User1 には次のロールがあります。

- \* 読者
- \* セキュリティ管理者
- \* セキュリティリーダー

User1 が VNet1 の閲覧者ロールを他のユーザーに割り当てることができるようにする必要があります。どうすればよいでしょうか？

- A.** User1 に VNet1 の共同作成者ロールを割り当てます。
- B.** Subscription1 のセキュリティ リーダーおよびリーダー ロールからユーザーを削除します。
- C.** User1 に VNet1 のネットワーク共同作成者ロールを割り当てます。
- D.** User1 に VNet1 のユーザー アクセス管理者ロールを割り当てます。

正解: **D** ([コメントを发表する](#))

<https://docs.microsoft.com/en-us/azure/role-based-access-control/rbac-and-directory-admin-roles#:~:>

text=The%20User%20Access%20Administrator%20role%20enables%20the%20user%20to%20grant,Azure%20subscriptions%20and%20management%20groups.

質問: **40**

アプリケーションを移動した後、App1 のバックアップ ソリューションを実装する必要があります。

最初に何を作成すればよいですか？

- A.** 回復計画
- B.** Azure バックアップ サーバー
- C.** バックアップポリシー
- D.** Recovery Services コンテナ

正解: ([正解を表示します](#))

A Recovery Services vault is a logical container that stores the backup data for each protected resource, such as Azure VMs. When the backup job for a protected resource runs, it creates a recovery point inside the Recovery Services vault.

Scenario:

There are three application tiers, each with five virtual machines.

Move all the virtual machines for App1 to Azure.

Ensure that all the virtual machines for App1 are protected by backups.

References: <https://docs.microsoft.com/en-us/azure/backup/quick-backup-vm-portal>


質問: 41

storage1 という名前の Azure ストレージ アカウントがあります。

Azure App Service アプリ「App1」と、Azure コンテナ インスタンスで実行されるアプリ「App2」があります。各アプリはマネージド ID を使用します。

今後 30 日間、App1 と App2 が storage1 から BLOB を読み取ることができることを確認する必要があります。

各アプリの storage1 で何を構成する必要がありますか？



正解:



Explanation:

Box 1: Access Control (IAM)

Since the App1 uses Managed Identity, App1 can access the Storage Account via IAM. As per requirement, we need to minimize the number of secrets used, so Access keys is not ideal.

Box 2: Shared access signatures (SAS)

We need temp access for App2, so we need to use SAS.

A shared access signature (SAS) provides secure delegated access to resources in your storage account without compromising the security of your data. With a SAS, you have granular control over how a client can access your data. You can control what resources the client may access, what permissions they have on those resources, and how long the SAS is valid, among other parameters.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-auth>

質問: 42

storageaccount1 という名前の Azure ストレージ アカウントを含む Azure サブスクリプションがあります。

storageaccount1 を Azure Resource Manager テンプレートとしてエクスポートします。このテンプレートには以下のセクションが含まれています。

```

{
  "type": "Microsoft.Storage/storageAccounts",
  "apiVersion": "2019-06-01",
  "name": "storageaccount1",
  "location": "eastus",
  "sku": {
    "name": "Standard_LRS",
    "tier": "Standard"
  },
  "kind": "StorageV2",
  "properties": {
    "networkAcls": {
      "bypass": "AzureServices",
      "virtualNetworkRules": [],
      "ipRules": [],
      "defaultAction": "Allow"
    },
    "supportsHttpsTrafficOnly": true,
    "encryption": {
      "services": {
        "file": {
          "keyType": "Account",
          "enabled": true
        },
        "blob": {
          "keyType": "Account",
          "enabled": true
        }
      }
    },
    "keySource": "Microsoft.Storage"
  },
  "accessTier": "Hot"
}

```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。  
 注意: 正しい選択ごとに1ポイントが加算されます。

Answer Area

Statements

Yes

No

A server that has a public IP address of 131.107.103.10 can access storageaccount1.

Individual blobs in storageaccount1 can be set to use the archive tier.

Global administrators in Azure AD can access a file share hosted in storageaccount1 by using their Azure AD credentials.

正解:

Answer Area

Statements

Yes

No

A server that has a public IP address of 131.107.103.10 can access storageaccount1.

Individual blobs in storageaccount1 can be set to use the archive tier.

Global administrators in Azure AD can access a file share hosted in storageaccount1 by using their Azure AD credentials.

Explanation:

Answer Area

Statements

Yes

No

A server that has a public IP address of 131.107.103.10 can access storageaccount1.

Individual blobs in storageaccount1 can be set to use the archive tier.

Global administrators in Azure AD can access a file share hosted in storageaccount1 by using their Azure AD credentials.

The provided ARM template defines a storage account named storageaccount1 of kind StorageV2 and SKU Standard\_LRS.

1## Public IP Access (Yes):

In the template, the networkAcls property sets " defaultAction " : " Allow " , meaning that access to the storage account is open to all public networks by default. Since no ipRules or virtualNetworkRules are defined, there are no restrictions based on IP addresses. Therefore, any server - including one with public IP

131.107.103.10 - can connect successfully.

(Reference: Azure Storage firewall and virtual network documentation - when defaultAction is Allow, storage accepts all incoming requests over the public endpoint.)

2## Archive Tier Availability (Yes):

Because the storage account is of kind StorageV2, it supports blob access tiers - Hot, Cool, and Archive - at both the account and individual blob levels. The ARM template specifies " accessTier " : " Hot " , which is the default access tier for the account. However, this does not restrict individual blobs from being set to a different tier (e.g., Archive) later. Azure Blob Storage in a StorageV2 account allows per-blob tiering, enabling you to optimize costs based on access patterns.

(Reference: Azure Storage documentation - " Blob storage lifecycle management " and " Blob access tiers overview. " )

3## Azure AD Authentication for File Shares (No):

The ARM template enables encryption and HTTPS but does not configure Azure Active Directory Domain Services (Azure AD DS) authentication for Azure Files. Global administrators in Azure AD cannot access SMB file shares using their Azure AD credentials unless Azure Files AD DS integration is explicitly configured. Since no such property (e.g., azureFilesIdentityBasedAuthentication) exists in this template, access would still require storage account keys or shared access signatures (SAS), not Azure AD credentials.

(Reference: Azure Files authentication overview - Azure AD DS authentication requires explicit configuration.)

# Final Verified Answers:

- \* A server that has public IP 131.107.103.10 can access storageaccount1 # Yes
- \* Individual blobs in storageaccount1 can be set to use the archive tier # Yes
- \* Global administrators in Azure AD can access a file share hosted in storageaccount1 by using their Azure AD credentials # No

質問: 43

Windows Server 2016 Datacenter イメージを使用する仮想マシン スケール セットの展開を自動化する予定です。  
スケール セット仮想マシンをプロビジョニングするときに、Web サーバー コンポーネントがインストールされていることを確認する必要があります。  
実行すべき 2 つのアクションはどれですか。それぞれの正解は解決策の一部を示しています。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

- A. Azure Resource Manager テンプレートの extensionProfile セクションを変更します。
- B. Azure ポータルで新しい仮想マシン スケール セットを作成します。
- C. Azure ポリシーを作成します。
- D. 自動化アカウントを作成します。
- E. 構成スクリプトをアップロードします。

正解: [\(正解を表示します\)](#)

To automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image and has web server components installed, you need to perform the following actions:  
Modify the extensionProfile section of the Azure Resource Manager template. This section defines the extensions that are applied to the scale set virtual machines after they are provisioned. You can use the Custom Script Extension to run PowerShell scripts that install and configure the web server components. For more information, see Deploy an application to an Azure Virtual Machine Scale Set1.  
Upload a configuration script. This is the PowerShell script that contains the commands to install and configure the web server components. You can upload the script to a storage account or a GitHub repository, and then reference it in the extensionProfile section of the template. For an example of a configuration script, see Tutorial: Install applications in Virtual Machine Scale Sets with Azure PowerShell2.

質問: 44

次の表に示す仮想ネットワークを含む Azure サブスクリプションがあります。

| Name  | Location   | Cloud type       |
|-------|------------|------------------|
| VNet1 | East US    | Azure Government |
| VNet2 | West US 2  | Public           |
| VNet3 | China East | Azure China      |

次の図に示すピアリング オプションがあります。



仮想ネットワーク上のリソースの通信戦略を設計する必要があります。  
以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

| Statements                               | Yes                   | No                    |
|------------------------------------------|-----------------------|-----------------------|
| Peering 1-2 is a possible configuration. | <input type="radio"/> | <input type="radio"/> |
| Peering 1-3 is a possible configuration. | <input type="radio"/> | <input type="radio"/> |
| Peering 3-2 is a possible configuration. | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                               | Yes                   | No                               |
|------------------------------------------|-----------------------|----------------------------------|
| Peering 1-2 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |
| Peering 1-3 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |
| Peering 3-2 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |

Explanation:

| Statements                               | Yes                   | No                               |
|------------------------------------------|-----------------------|----------------------------------|
| Peering 1-2 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |
| Peering 1-3 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |
| Peering 3-2 is a possible configuration. | <input type="radio"/> | <input checked="" type="radio"/> |

Azure Virtual Network (VNet) peering allows two virtual networks to communicate privately across the Azure backbone network. According to Microsoft Azure Administrator documentation, VNet peering can be of two main types:

Regional VNet Peering: within the same Azure region and cloud type.

Global VNet Peering: across Azure regions, but within the same Azure cloud.

Microsoft Azure operates multiple sovereign clouds, which are logically and physically isolated:

Azure Global/Public Cloud (used by most commercial customers).

Azure Government Cloud (for U.S. government agencies and contractors, physically isolated from the public cloud).

Azure China Cloud (operated by 21Vianet, completely independent of Microsoft's global Azure).

Per Microsoft Learn - Virtual Network Peering Documentation, the following restrictions apply:

"VNet peering cannot be established between Azure sovereign clouds (such as between Azure Government, Azure China, and Azure Public regions). Peering is only supported within the same cloud type, whether within a region or across regions." Based on the table:

VNet1 (East US, Azure Government)

VNet2 (West US 2, Public)

VNet3 (China East, Azure China)

Analysis of each peering configuration:

Peering 1-2 (Azure Government # Public Cloud): # Not supported.

Azure Government and Azure Public are separate environments - no peering across these cloud boundaries.

Peering 1-3 (Azure Government # Azure China): # Not supported.

These are two isolated sovereign clouds with independent infrastructure.

Peering 3-2 (Azure China # Public Cloud): # Not supported.

Azure China is operated by 21Vianet and cannot connect directly with the Public Azure cloud.

Therefore, no cross-cloud peering is supported among these networks.

If all VNets were within Public Azure (e.g., East US, West US 2, and North Europe), global VNet peering would be allowed. However, in this scenario, each VNet resides in a different cloud type, so no peering is possible between them.

質問: **45**

VM1という名前のAzure仮想マシンをバックアップする予定です。

バックアップ事前チェックのステータスが「警告」と表示されていることに気づきます。

警告ステータスが表示される原因として考えられるものは何ですか？

**A.** リカバリサービス ボールトが利用できません。

**B.** VM1 には管理されていないディスクがあります。

**C.** VM1 には最新バージョンの Azure VM エージェント (WaAppAgent.exe) がインストールされていません。

**D.** VM1 が停止しました。

正解: ([正解を表示します](#))

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

Azure Backup Pre-Check validates whether the VM configuration is suitable for successful backup before or during backup processing. In Microsoft's description of Azure VM backup pre-checks, a Warning state indicates "one or more issues in VM's configuration that might lead to backup failures," and Microsoft specifically gives the example that not having the latest VM Agent installed can cause intermittent backup failures and falls into this Warning category.

Therefore, the correct answer is C. The Azure VM Agent, shown on Windows VMs as WaAppAgent.exe, enables communication between the Azure platform, VM extensions, and backup components. Azure Backup relies on this integration to install and manage the VM snapshot extension and coordinate backup operations.

Microsoft troubleshooting guidance also lists ensuring that the VM Agent is the latest version as a core backup troubleshooting requirement.

The other options do not best match a Warning pre-check status. A vault availability issue is not a VM configuration warning. An unmanaged disk is not the expected pre-check warning cause here. A stopped VM state alone is not the cited reason for this warning. Study Guide reference: AZ-104 "Monitor and back up Azure resources" > "Implement backup and recovery" > Azure Backup / Recovery Services vault / Azure VM backup pre-check status.

質問: **46**

8 台の仮想マシンと次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name     | Description                                                                    |
|----------|--------------------------------------------------------------------------------|
| storage1 | Storage account                                                                |
| storage2 | Storage account                                                                |
| VNET1    | Virtual network with a single subnet that has five virtual machines connected  |
| VNET2    | Virtual network with a single subnet that has three virtual machines connected |

VNET1 へのアクセスを設定する必要があります。ソリューションは以下の要件を満たしている必要があります。

- \* VNET1 に接続された仮想マシンは、Microsoft バックボーンを使用して VNET2 に接続された仮想マシンと通信できる必要があります。
- \* VNET1 に接続された仮想マシンは、Microsoft バックボーンを使用して、storage1、storage2、および Microsoft Entra ID にアクセスできる必要があります。

VNET1 に追加する必要があるサービス エンドポイントの最小数はいくつですか？

- A. 1
- B. 2
- C. 3
- D. 5

正解: (正解を表示します)

When you perform a VM restore using Azure Backup with the "Replace existing" option:

- \* Azure Backup replaces the existing VM's operating system disk and configuration with what was captured at the time of the backup.
- \* Data disks added after the backup are not part of the recovery point and are therefore not restored.

Other post-backup changes:

- \* VM size change: Automatically reverted to the size captured in the backup configuration.
- \* File addition (Budget.xls): Not present, because only data up to the backup time is restored.
- \* Password reset: Resetting password does not persist because OS disk from backup overwrites current one.However, when restoring, you can always reset credentials later - but that is expected behavior.

The only configuration that requires manual re-creation after restore is any new data disk added after the backup snapshot.

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 47

VM1とVM2をLB1のバックエンドに追加する必要があります。まず何をすべきでしょうか？

- A. 新しい NSG を作成し、NSG を VNET1/Subnet1 に関連付けます。
- B. VM2 を VNET1/Subnet1 に接続します。
- C. VM1 と VM2 を同じ可用性ゾーンに再デプロイします。
- D. VM1 と VM2 を同じ可用性セットに再デプロイします。

正解: (正解を表示します)

In Azure, Load Balancers distribute network traffic across multiple virtual machines (VMs) to ensure high availability and reliability. To add virtual machines to the backend pool of an Azure Load Balancer, the following key conditions must be met according to the Microsoft Azure Administrator documentation:

All VMs in the backend pool must be connected to the same virtual network (VNet) as the Load Balancer.

The Load Balancer (in this case, LB1) is configured for internal load balancing on VNET1/Subnet1 as per the technical requirements of the case study.

The backend pool can include network interfaces (NICs) from VMs within the same region and VNet.

Step-by-step analysis:

From the case study data:

VM

Location

Connected to

IP Address

VM1

West US

VNET1/Subnet1

10.0.1.4

VM2

West US

VNET1/Subnet2

10.0.2.4

LB1

Internal Basic Load Balancer

Connected to VNET1/Subnet1

-

Observation:

VM1 is already connected to VNET1/Subnet1, where the internal Load Balancer LB1 is also deployed.

VM2, however, is connected to VNET1/Subnet2, which is a different subnet within the same virtual network.

According to Microsoft Learn ("Configure backend pools in Azure Load Balancer"):

"All network interfaces in the backend pool must be within the same virtual network as the load balancer.

You cannot add VMs connected to different VNets or subnets not associated with the load balancer's front- end configuration." Therefore, before you can add VM2 to the backend pool, you must ensure that its network interface is attached to VNET1/Subnet1, the same subnet used by LB1.

Only after this step will both VMs (VM1 and VM2) be eligible for inclusion in LB1's backend pool.

Incorrect Option Analysis:

A). Create a new NSG and associate the NSG to VNET1/Subnet1.

# Not required. Network Security Groups control traffic filtering, not backend pool configuration.

C). Redeploy VM1 and VM2 to the same availability zone.

# Availability Zones only matter for redundancy and failover, not for backend pool eligibility in a basic internal load balancer.

D). Redeploy VM1 and VM2 to the same availability set.

# Basic Load Balancers can distribute traffic across VMs in the same availability set, but both VMs must already reside in the same VNet/Subnet first.

Final Verified Answer:

# B. Connect VM2 to VNET1/Subnet1

Reference (Microsoft Official Documentation):

Microsoft Learn: Configure the backend pool for Azure Load Balancer

Microsoft Learn: Azure Load Balancer overview

Microsoft Learn: Create and configure an internal load balancer

Microsoft Learn: Virtual network and subnet requirements for load balancing

質問: 48

次の表に示すリソースが含まれる、Subscription! という名前の Azure サブスクリプションがあります。

| Name     | Type                    | Region     | Resource group |
|----------|-------------------------|------------|----------------|
| RG1      | Resource group          | West US    | Not applicable |
| RG2      | Resource group          | West US    | Not applicable |
| Vault1   | Recovery Services vault | Central US | RG1            |
| Vault2   | Recovery Services vault | West US    | RG2            |
| VM1      | Virtual machine         | Central US | RG2            |
| storage1 | Storage account         | West US    | RG1            |
| SQL1     | Azure SQL database      | East US    | RG2            |

storage1 に、blob1 という名前の BLOB コンテナと share1 という名前のファイル共有を作成します。

Vault1 と Vault2 にバックアップできるリソースはどれですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Can use Vault1 for backups:

VM1 only

VM1 only

VM1 and share1 only

VM1 and SQL1 only

VM1, storage1, and SQL1 only

VM1, blob1, share1, and SQL1

Can use Vault2 for backups:

share1 only

storage1 only

share1 only

VM1 and share1 only

blob1 and share1 only

storage1 and SQL1 only

正解:

Answer Area

Can use Vault1 for backups:

VM1 only

VM1 only

VM1 and share1 only

VM1 and SQL1 only

VM1, storage1, and SQL1 only

VM1, blob1, share1, and SQL1

Can use Vault2 for backups:

share1 only

storage1 only

share1 only

VM1 and share1 only

blob1 and share1 only

storage1 and SQL1 only

Explanation:

Answer Area

Can use Vault1 for backups:

VM1 only

Can use Vault2 for backups:

share1 only

Azure Recovery Services vaults are used to store backup data and recovery points for supported workloads such as Azure VMs, Azure Files, and on-premises workloads. However, the configuration rules are strictly defined based on region and supported resource type.

From Microsoft Azure Administrator documentation and the official Azure Backup architecture, the following principles apply:

\* Backup Region Dependency - A Recovery Services vault can only back up resources located in the same Azure region. For example, a vault in Central US can only back up Azure VMs, storage accounts, or workloads in Central US. It cannot back up resources in West US or East US.

\* Vault1 (Central US, RG1)

\* Located in Central US and can therefore back up only resources in Central US.

\* From the table: VM1 is in Central US, and is eligible.

\* storage1 (West US) and SQL1 (East US) cannot be backed up to Vault1 due to cross-region restrictions.

# Therefore, Vault1 can back up VM1 only.

\* Vault2 (West US, RG2)

\* Located in West US, and can only back up resources in West US.

\* In West US, the only relevant resource is storage1, which contains a file share named share1.

Azure Backup supports Azure File share backup via Recovery Services vaults.

\* Vault2 cannot back up VM1 (Central US) or SQL1 (East US) since they reside in different regions.

# Therefore, Vault2 can back up share1 only.

\* Azure Backup Limitations Summary (From Microsoft Docs):

\* Azure VMs # must be backed up to a vault in the same region.

\* Azure File Shares # can be backed up using Recovery Services vaults in the same region as the storage account.

\* Azure SQL Database # uses its own automated backup mechanism, not Recovery Services vaults.

Hence, by Azure documentation compliance and AZ-104 study guide principles:

\* Vault1 backs up VM1 only.

\* Vault2 backs up share1 only.

Final Verified Answer:

# Vault1 - VM1 only

# Vault2 - share1 only

質問: 49


注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、指定された目標を達成できる独自のソリューションが含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答すると、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

VM1とVM2という名前の2つのAzure仮想マシンにインストールされているApp1という名前のアプリがあります。

App1への接続は、Azureロードバランサーを使用して管理されます。

VM2の効果的なネットワークセキュリティ構成を次の図に示します。



TCPポート443を介した131.107.100.50からApp1への接続が失敗することがわかります。ロードバランサのルールが正しく構成されていることを確認します。

App1への接続が、TCPポート443経由で131.107.100.50から正常に確立できることを確認する必要があります。

解決策 :Allow\_131.107.100.50インバウンドセキュリティルールの優先度を変更します。

これは目標を達成していますか？

A. はい

B. いいえ

正解: A (コメントを发表する)

In this scenario, you are troubleshooting inbound network traffic to VM2 where connections from the public IP 131.107.100.50 over TCP port 443 are failing, despite the Azure Load Balancer configuration being correct. The issue lies in the Network Security Group (NSG) rules applied to VM2-NIC1.

Understanding NSG Rule Processing

According to Microsoft Azure Administrator documentation (AZ-104 Study Guide, "Secure network traffic" section):

"Network security group rules are processed in priority order, from the lowest number to the highest number.

As soon as a rule matches the traffic, rule processing stops."

This means that once a matching rule is found (based on source, destination, port, and protocol), Azure will not evaluate further rules.

Analysis of the Exhibit

From the NSG configuration shown:

| Priority | Name                          | Port | Protocol | Source         | Destination       | Action |
|----------|-------------------------------|------|----------|----------------|-------------------|--------|
| 100      | Allow_131.107.100.50          | 443  | TCP      | 131.107.100.50 | VirtualNetwork    | Allow  |
| 200      | Block_All_Other443            | 443  | TCP      | Any            | Any               | Deny   |
| 65000    | AllowVNetInBound              |      |          | Any            | Any               | Allow  |
| 65001    | AllowAzureLoadBalancerInBound |      |          | Any            | AzureLoadBalancer | Allow  |
| 65500    | DenyAllInBound                |      |          | Any            | Any               | Deny   |

Any

Deny

At first glance, Rule 100 should allow traffic from 131.107.100.50 on port 443, but it specifies Destination:

VirtualNetwork, which applies only if the destination IP is part of the same virtual network (VNet).

However, in this case, the traffic originates from an external public IP address (131.107.100.50) trying to reach VM2 through the Azure Load Balancer, not directly through the virtual network.

Because of this, the destination does not match the "VirtualNetwork" scope, and the rule is effectively skipped. The next applicable rule (priority 200) explicitly denies all traffic on port 443. Therefore, the connection fails.

Proposed Solution: Modify the Priority

The question's solution proposes modifying the priority of the Allow\_131.107.100.50 rule.

If you adjust the rule's priority and ensure that it is processed before the "Block\_All\_Other443" deny rule (priority 200) and correct its destination to Any, then the traffic from 131.107.100.50 will be evaluated first and allowed before the deny rule triggers.

Hence, by modifying the priority or adjusting the target scope, the traffic will successfully match the allow rule.

Official Documentation Reference

From Microsoft Learn - Network Security Groups Overview:

"Rules are processed in priority order. Once a rule matches the traffic, processing stops."

"If you have a rule that denies specific traffic with a higher priority, any allow rules with a lower priority won't be processed." Conclusion By modifying the priority of the Allow\_131.107.100.50 rule (or ensuring it appears before the deny rule in evaluation order), inbound traffic from that specific IP over port 443 will be allowed before the deny rule applies.

Therefore, the proposed solution meets the goal.

# Final Verified Answer: A. Yes

Modifying the priority of the "Allow\_131.107.100.50" rule ensures it is processed before the

"Block\_All\_Other443" deny rule, allowing the required inbound traffic to reach the VM successfully.

質問: 50

VM1 という名前の Azure 仮想マシンがあります。VM1 は、ARMtjson という名前のカスタム Azure Resource Manager テンプレートを使用してデプロイされました。

VM1 がメンテナンスの影響を受けるという通知を受け取ります。

VM1 をすぐに別のホストに移動する必要があります。

解決策: VM1 更新ブレードから、1 回限りの更新を選択します。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: (正解を表示します)

When a virtual machine (VM) in Azure is scheduled for platform maintenance, it can either be planned maintenance (with notification) or unplanned maintenance (without notice). During planned maintenance, administrators have the option to move the VM to a different host proactively to minimize downtime.

According to Microsoft Azure documentation ("Maintenance and updates for Azure virtual machines"), the correct action is to use the "Redeploy" or "Move to another host" feature, not a "One-time update." The "Redeploy" operation deallocates the VM and migrates it to a new node within Azure's infrastructure, resolving underlying host issues while preserving OS and data disks.

The "Updates blade" in the Azure portal does not contain a "One-time update" option to move a VM off a host. That term is used in the context of patching or guest OS updates, not host maintenance.

Therefore, the proposed solution does not meet the goal of moving the VM to another host immediately. The correct action would be to redeploy the VM using the Azure portal or Invoke-AzVMRedeploy PowerShell cmdlet.

質問: 51

Subscription1 と Subscription2 という 2 つのサブスクリプションがあります。各サブスクリプションは異なる Azure AD テナントに関連付けられています。

Subscription1 には、VNet1 という名前の仮想ネットワークが含まれています。VNet1 には、VM1 という名前の Azure 仮想マシンが含まれており、IP アドレス空間は 10.0.0.0/16 です。

Subscription2 には、VNet2 という名前の仮想ネットワークが含まれています。VNet2 には、VM2 という名前の Azure 仮想マシンが含まれており、IP アドレス空間は 10.10.0.0/24 です。

VNet1 を VNet2 に接続する必要があります。

まず何をすべきでしょうか？

**A.** VM1 をサブスクリプション 2 に移動します。

**B.** VNet2 の IP アドレス空間を変更します。

**C.** 仮想ネットワーク ゲートウェイをプロビジョニングします。

**D.** VNet1 を Subscription2 に移動します。

正解: ([正解を表示します](#))

Azure VNet peering is the most common way to connect virtual networks, but in this scenario, VNet1 and VNet2 exist in different Azure AD tenants. Global or cross-tenant peering requires that both VNets be within the same Azure Active Directory tenant.

Because the question specifies that each subscription is associated with a different Azure AD tenant, VNet peering is not supported across tenants.

To establish connectivity between VNets in different tenants, the correct solution is to configure VPN gateways in both virtual networks and create a VNet-to-VNet VPN connection.

This type of connection uses the Azure VPN Gateway service to securely tunnel traffic between the VNets using IPsec/IKE protocols over the Internet.

Changing IP address spaces or moving VNets/subscriptions does not solve the cross-tenant issue securely or efficiently.

Hence, before connecting VNet1 and VNet2, you must provision virtual network gateways in both VNets.

質問: 52

10 個の Azure サブスクリプションにリンクされている Azure AD テナントがあります。すべてのサブスクリプションでユーザー アクティビティを一元的に監視する必要があります。何を使うべきですか？

**A.** アクティビティ ログ フィルター

**B.** Log Analytics ワークスペース

**C.** アクセスレビュー

**D.** Azure Application Insights プロファイラー

正解: ([正解を表示します](#))

<https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/activity-log?tabs=powershell#send-to-log-analytics-workspace> Send the activity log to a Log Analytics workspace to enable the Azure Monitor Logs feature, where you: - Consolidate log entries from multiple Azure subscriptions and tenants into one location for analysis together.

質問: 53

5,000 個のユーザー アカウントを含む Microsoft Entra テナントがあります。

AdminUser1 という名前の新しいユーザー アカウントを作成します。

AdminUser1 にユーザー管理者の管理ロールを割り当てする必要があります。

ユーザー アカウントのプロパティから何をすればよいですか？

**A.** [グループ] ブレードから、ユーザー アカウントを新しいグループに招待します。

**B.** ディレクトリ ロール ブレードから、ディレクトリ ロールを変更します。

**C.** [ライセンス] ブレードから、新しいライセンスを割り当てます。

正解: ([正解を表示します](#))

In Microsoft Entra ID (formerly Azure Active Directory), roles are assigned to users to delegate administrative permissions in a least-privilege manner. The User Administrator role allows a user to manage other users and groups - for example, creating and managing user accounts, resetting passwords for non-administrators, and managing user group memberships.

To assign a role such as User Administrator, you must use the Directory role blade within the user's account properties in the Azure portal.

Step-by-step according to Microsoft documentation:

\* Sign in to the Azure Portal using an account that has one of the following roles:

- \* Global Administrator
- \* Privileged Role Administrator
- \* Navigate to Azure Active Directory # Users # select AdminUser1.
- \* Under Manage, select Directory role.This blade shows all current role assignments for the selected user.
- \* Click Add assignment (or Modify role).
- \* Select the User Administrator role from the list of available directory roles, then click Add.

Once this is completed, AdminUser1 will have administrative permissions limited to user management activities within the tenant.

Why other options are incorrect:

- \* A. From the Groups blade, invite the user account to a new group:Group membership does not grant directory-level administrative permissions. Roles must be assigned at the directory role level, not via groups (unless using role-assignable groups configured for PIM).
- \* C. From the Licenses blade, assign a new license:Licenses determine service usage (e.g., Microsoft 365, Intune) and do not provide administrative privileges in Entra ID.

Extract from Microsoft Azure Administrator Documentation (Official Guide):

"To assign a role to a user, in the Azure portal, select the user, then under Manage select Directory role, and choose the role you want to assign." (Source: Microsoft Learn - Assign roles to users in Azure Active Directory)

質問: 54

VNet1 と VNet2 という 2 つのピアリングされた仮想ネットワークを含む Azure サブスクリプションがあります。VNet1 には、静的ルーティングを使用する VPN ゲートウェイがあります。

オンプレミス ネットワークには、VNet1 の VPN ゲートウェイを使用する VPN 接続があります。

オンプレミス ネットワーク上のユーザーが VNet2 上の仮想マシンに接続できるようにアクセスを構成する必要があります。

ソリューションではコストを最小限に抑える必要があります。

どのようなタイプの接続を使用する必要がありますか？

- A. プライベート IP アドレスを持つ Azure Firewall
- B. VNet2へのExpressRoute回線
- C. サービスチェイニングとユーザー定義ルート (UDR)
- D. Azure アプリケーション ゲートウェイ

正解: C ([コメントを发表する](#))

Because VNet1 and VNet2 are peered, and VNet1 already has a VPN gateway using static routing, the most cost-effective method is to use service chaining with UDRs.

This allows:

- \* Traffic from on-premises to traverse the VPN gateway
- \* Routing across peered VNets without deploying additional gateways
- \* No additional ExpressRoute or firewall costs

Microsoft documentation states:

"User-defined routes enable custom routing across peered virtual networks for transit scenarios." Azure Firewall and Application Gateway add unnecessary cost. ExpressRoute is the most expensive option.

質問: 55

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、指定された目標を達成する可能性のある独自のソリューションが含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これら

質問はレビュー画面に表示されません。

West US AzureリージョンでホストされているVNet1という名前の仮想ネットワークを管理します。

VNet1は、Windows Serverを実行するVM1およびVM2という名前の2つの仮想マシンをホストします。

VM1からVM2へのすべてのネットワークトラフィックを3時間検査する必要があります。

解決策 パフォーマンスモニターから、データコレクターセット (DCS)を作成します。

これは目標を達成していますか？

A. はい

B. いいえ

正解: B (コメントを发表する)

To inspect or capture network traffic between two Azure virtual machines (VM1 and VM2) within the same virtual network, Azure provides specialized network diagnostic tools - Azure Network Watcher, Connection Monitor, and Packet Capture - rather than Windows Performance Monitor.

The Performance Monitor (PerfMon) tool on Windows is used to collect metrics such as CPU, memory, and disk I/O performance counters, not network packet data or traffic analysis between VMs. While you could use PerfMon to track bandwidth usage or network throughput locally, it cannot capture or inspect packet-level communication or determine the content or flow between two Azure VMs.

According to the Microsoft Azure Administrator documentation (AZ-104 study guide), the correct method to analyze and inspect network traffic between Azure virtual machines includes:

Enabling Azure Network Watcher in the target region.

Using Packet Capture under Network Watcher to capture inbound and outbound packets.

Optionally using Connection Monitor to monitor connectivity metrics between endpoints.

Packet Capture allows you to define capture filters (e.g., source/destination IP, ports, protocols) and run it for a specific duration, such as three hours, as required in this scenario. The capture file is stored in Azure Storage or locally for later analysis with tools like Wireshark.

Therefore, creating a Data Collector Set (DCS) in Performance Monitor does not meet the requirement to inspect network traffic between VM1 and VM2.

質問: 56

次の表に示すリソースが含まれる Sub1 という名前の Azure サブスクリプションがあります。

| Name    | Description                                                    |
|---------|----------------------------------------------------------------|
| RG1     | Resource group                                                 |
| Action1 | Action group that sends an email message to admin1@contoso.com |

Sub1 には次のアラート ルールが含まれています。

\* 名前: アラート1

\* スコープ: Sub1 内のすべてのリソース グループ

o 将来のすべてのリソースを含める

\* 条件: すべての管理操作

\* アクション: アクション1

Sub1 には次のアラート処理ルールが含まれています。

\* 名前: ルール1

\* 範囲: サブ1

\* ルールの種類: 通知を抑制する

\* ルールを適用: 特定の時間に

テントについて: 2022年8月10日

終了日: 2022年8月13日

以下の各文について、正しい場合は 「はい」を選択してください。そうでない場合は 「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

| Answer Area                                                                                                |  |  | Yes                   | No                    |
|------------------------------------------------------------------------------------------------------------|--|--|-----------------------|-----------------------|
| Statements                                                                                                 |  |  |                       |                       |
| If you create a resource group in Sub1 on August 11, 2022, Alert1 is listed in the Azure portal.           |  |  | <input type="radio"/> | <input type="radio"/> |
| If you create a resource group in Sub1 on August 12, 2022, an email message is sent to admin1@contoso.com. |  |  | <input type="radio"/> | <input type="radio"/> |
| If you add a tag to RG1 on August 15, 2022, an email message is sent to admin1@contoso.com.                |  |  | <input type="radio"/> | <input type="radio"/> |

正解:

| ANSWER AREA                                                                                                |  |  | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------|--|--|----------------------------------|----------------------------------|
| Statements                                                                                                 |  |  |                                  |                                  |
| If you create a resource group in Sub1 on August 11, 2022, Alert1 is listed in the Azure portal.           |  |  | <input checked="" type="radio"/> | <input type="radio"/>            |
| If you create a resource group in Sub1 on August 12, 2022, an email message is sent to admin1@contoso.com. |  |  | <input type="radio"/>            | <input checked="" type="radio"/> |
| If you add a tag to RG1 on August 15, 2022, an email message is sent to admin1@contoso.com.                |  |  | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

| Answer Area                                                                                                |  |  | Yes                              | No                               |
|------------------------------------------------------------------------------------------------------------|--|--|----------------------------------|----------------------------------|
| Statements                                                                                                 |  |  |                                  |                                  |
| If you create a resource group in Sub1 on August 11, 2022, Alert1 is listed in the Azure portal.           |  |  | <input checked="" type="radio"/> | <input type="radio"/>            |
| If you create a resource group in Sub1 on August 12, 2022, an email message is sent to admin1@contoso.com. |  |  | <input type="radio"/>            | <input checked="" type="radio"/> |
| If you add a tag to RG1 on August 15, 2022, an email message is sent to admin1@contoso.com.                |  |  | <input checked="" type="radio"/> | <input type="radio"/>            |

This question evaluates your understanding of Azure Monitor alerts, alert processing rules, and action groups.

1. Alert1 Configuration Details

Scope: All resource groups in Sub1

Include all future resources: # Yes

Condition: All administrative operations

This means any action that modifies Azure resources (create, delete, update, tagging, etc.) triggers the alert.

Action Group: Action1 # Sends email to admin1@contoso.com

So, Alert1 applies to all current and future resource groups in the subscription and triggers when any administrative operation occurs within that scope.

2. Alert Processing Rule (Rule1)

Scope: Sub1 (entire subscription)

Rule type: Suppress notifications

Start: August 10, 2022

End: August 13, 2022

According to Microsoft documentation:

"When a suppress notification rule is active, alert notifications (such as emails or SMS) are not sent, but the alerts themselves are still created and listed in Azure Monitor." So, between August 10-13, 2022, alerts will still appear in the Azure portal, but no email notifications are sent.

Now evaluate each statement:

Statement 1:

"If you create a resource group in Sub1 on August 11, 2022, Alert1 is listed in the Azure portal." The action (creating a resource group) triggers Alert1, since it's an administrative operation within scope.

However, August 11 is during the suppression period.

The alert is still created and visible in the Azure portal, but the notification is suppressed.

# Answer: Yes (Alert appears in the portal even during suppression)

Statement 2:

"If you create a resource group in Sub1 on August 12, 2022, an email message is sent to admin1@contoso.com."

Date: August 12, 2022 - still within the suppression period (Aug 10-13).

Therefore, no email notification is sent.

# Answer: No

Statement 3:

"If you add a tag to RG1 on August 15, 2022, an email message is sent to admin1@contoso.com." Date: August 15, 2022, which is after the suppression period (ended Aug 13).

The alert condition ("All administrative operations") includes tag changes.

Therefore, this operation triggers an alert and sends an email through Action1.

# Answer: Yes

Official Microsoft Azure Documentation Extract (AZ-104 Study Material Reference):

"Alert processing rules can temporarily suppress or redirect alert notifications without disabling or deleting the alert rule itself."

"When a suppress notification rule is active, alerts are still created in the Azure Monitor activity log, but notifications (emails, SMS, webhooks) are not sent."

"Administrative operations include all create, update, and delete actions at the resource and resource group level."

質問: 57

storage1 という名前のAzure Storageアカウントがあります。

User1 という名前のユーザーが storage1 のストレージ アカウント キーを一覧表示および再生成できるようにする必要があります。

解決策: ストレージ アカウント暗号化スコープの共同作成者ロールを User1 に割り当てます。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: ([正解を表示します](#))

In Azure, storage account keys are managed through Azure Role-Based Access Control (RBAC) permissions.

The ability to list and regenerate storage account keys requires access to specific management operations on the storage account resource, not encryption scopes.

According to Microsoft Azure Administrator documentation, the permissions required to list and regenerate storage account keys are:

\* Microsoft.Storage/storageAccounts/listkeys/action

\* Microsoft.Storage/storageAccounts/regeneratekey/action

Only specific built-in roles include these actions, such as:

\* Owner

\* Contributor

\* Storage Account Contributor

However, the role mentioned in the question - Storage Account Encryption Scope Contributor - is designed for a different purpose.

From Microsoft's official documentation for Storage Account Encryption Scope Contributor:

"Grants permissions to manage encryption scopes within a storage account. Does not grant access to manage storage account keys or data." This means that assigning the Storage Account Encryption Scope Contributor role gives the user permissions to configure encryption scopes (for example, enabling or modifying encryption at the blob container level), but not to list or regenerate access keys.

Therefore, the solution fails to meet the requirement of enabling User1 to list and regenerate keys.

To achieve the goal, the correct role would be:

# Storage Account Contributor

This role explicitly allows managing all storage account properties, including key management.

Key Point (Microsoft Official Documentation Summary):

- \* Storage Account Encryption Scope Contributor # Manage encryption scopes only.
- \* Storage Account Contributor # Manage storage account settings, including listing and regenerating keys.

質問: 58

VM1 という名前の仮想マシンを含む Azure サブスクリプションがあります。

DC1 というドメイン コントローラーを含むオンプレミス データセンターがあります。オンプレミス データセンターを Azure に接続するために ExpressRoute が使用されます。

VM1 と DC1 間のネットワーク遅延を識別するには、接続モニターを使用する必要があります。

DC1 に何をインストールすればよいですか？

- A. Log Analytics エージェント
- B. Azure Network Watcher Agent 仮想マシン拡張機能
- C. Azure Monitor エージェント拡張機能
- D. Azure Arc 対応サーバー用の Azure Connected Machine エージェント

正解: [\(正解を表示します\)](#)

This question focuses on how to monitor network connectivity and latency between Azure virtual machines and on-premises resources using Azure Network Watcher - Connection Monitor (v2).

# Scenario Breakdown

VM1: Azure virtual machine (in your subscription)

DC1: On-premises domain controller (in your datacenter)

Connectivity: Via ExpressRoute

Goal: Use Connection Monitor to track network latency between VM1 (Azure) and DC1 (on-premises) To achieve this, both endpoints (VM1 and DC1) must have agents capable of collecting and sending network telemetry data to Azure Monitor.

# Understanding Azure Connection Monitor (v2)

According to Microsoft Learn ("Monitor network connectivity with Connection Monitor"):

"Connection Monitor (v2) enables you to monitor network connectivity between Azure and on-premises resources. You can monitor connections between Azure VMs, Azure Arc-enabled servers, and any endpoint reachable over TCP or ICMP." To participate in a hybrid connection, the on-premises machine must be onboarded to Azure Arc.

Azure Arc connects your non-Azure servers to Azure Resource Manager and allows management and monitoring using Azure services, including Network Watcher's Connection Monitor.

# Required Agent for On-Premises Monitoring

For on-premises servers (like DC1), Azure Arc uses the Azure Connected Machine agent (formerly known as the Azure Arc agent).

This agent:

Registers the on-premises machine as an Azure Arc-enabled server in Azure.

Enables the use of Azure Monitor, Defender for Cloud, Update Management, and Connection Monitor on that server.

Once the Azure Connected Machine agent is installed and the server is connected through Azure Arc, you can add it as a source or destination endpoint in Connection Monitor to measure latency and packet loss.

# Why Other Options Are Incorrect

Option

Description

Why Incorrect

A). Log Analytics agent

Used for data collection (logs and metrics) for Azure Monitor.

# Does not support Connection Monitor (v2) endpoint monitoring. Deprecated in favor of Azure Monitor Agent.

B). Azure Network Watcher Agent extension

Used only on Azure VMs, not on on-premises servers.

# Cannot be installed on DC1 (non-Azure).

C). Azure Monitor agent extension

Used for telemetry/log ingestion into Azure Monitor.

# Does not support connectivity monitoring for hybrid endpoints.

# D. Azure Connected Machine agent

Connects on-premises servers to Azure Arc, enabling Connection Monitor and other Azure services.

# Correct and required.

# Verification (Microsoft Documentation Extract)

From Microsoft Learn - "Monitor hybrid connectivity using Connection Monitor":

"To monitor on-premises resources, onboard the servers to Azure Arc and install the Azure Connected Machine agent. This agent enables Connection Monitor to collect network connectivity data from on-premises endpoints."

# Final Verified Answer:

D). the Azure Connected Machine agent for Azure Arc-enabled servers

Summary of Key Points

Connection Monitor (v2) can track network latency between Azure and on-premises systems.

On-premises servers must be Azure Arc-enabled using the Azure Connected Machine agent.

Azure VMs use the Network Watcher extension, while non-Azure machines require Azure Arc for integration.

# Correct Answer: D. the Azure Connected Machine agent for Azure Arc-enabled servers

質問: 59

注: この質問は、同じシナリオを提示する一連の質問の一部です。一連の質問にはそれぞれ、定められた目標を満たす可能性のある独自の解答が含まれています。質問セットによっては、複数の正解が存在する場合もあれば、正解がない場合もあります。

このセクションの質問に回答した後は、その質問に戻ることはできません。そのため、これらの質問はレビュー画面に表示されません。

Adatum.com という名前の Microsoft Entra テナントと、Subscription1 という名前の Azure サブスクリプションがあります。

Adatum.com には Developers というグループが含まれています。Subscription1 には Dev というリソース グループが含まれています。

開発者グループに、Dev リソース グループに Azure Logic Apps を作成する権限を付与する必要があります。

解決策: Subscription1 で、DevTest Labs ユーザー ロールを開発者グループに割り当てます。

これは目標を満たしていますか？

**A.** はい

**B.** いいえ

正解: ([正解を表示します](#))

The DevTest Labs User role is a built-in Azure role designed specifically to support operations within Azure DevTest Labs. According to the Microsoft Azure Administrator documentation on built-in roles, this role permits users to connect to, start, stop, restart, and use virtual machines inside a DevTest Lab, but it does not grant permissions to create or manage Azure resources outside of DevTest Labs, such as Azure Logic Apps.

In this scenario, the requirement is to allow the Developers group to create Azure Logic Apps in the Dev resource group. Azure Logic Apps are managed through the Microsoft.Logic resource provider and require permissions such as Microsoft.Logic/workflows/write. These permissions are included in roles like Logic App Contributor or the more general Contributor role at the resource group scope.

The DevTest Labs User role does not include permissions to deploy or manage Logic Apps, even when assigned at the subscription level. Therefore, assigning this role does not meet the stated goal. Microsoft documentation explicitly states that DevTest Labs User is limited to DevTest Lab resources and cannot be used as a general deployment role.

Final Answer: B. No

質問: 60

Azure サブスクリプションをお持ちです。

次の Azure Resource Manager (ARM) テンプレートを使用して、storage1 という名前のストレージ アカウントをデプロイする予定です。

```
{
  "$schema": "http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "resources": [
    {
      "name": "storage1",
      "type": "Microsoft.Storage/storageAccounts",
      "apiVersion": "2021-08-01",
      "location": "East US",
      "properties": {
        "allowBlobPublicAccess": true,
        "defaultToOAuthAuthentication": false,
        "networkAcls": {
          "bypass": "AzureServices",
          "defaultAction": "Allow",
          "ipRules": []
        },
        "isVersioningEnabled": true
      },
      "dependsOn": [
        "[resourceId('Microsoft.Storage/storageAccounts/', 'storage1')]"
      ]
    }
  ]
}
```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

|                                                                           |                       |                       |
|---------------------------------------------------------------------------|-----------------------|-----------------------|
| Changes made to the data in storage1 can be rolled back after seven days. | <input type="radio"/> | <input type="radio"/> |
| Only users located in the East US Azure region can connect to storage1.   | <input type="radio"/> | <input type="radio"/> |
| Three copies of storage1 will be maintained in the East US Azure region.  | <input type="radio"/> | <input type="radio"/> |

正解:

|                                                                           |                                  |                                  |
|---------------------------------------------------------------------------|----------------------------------|----------------------------------|
| Changes made to the data in storage1 can be rolled back after seven days. | <input checked="" type="radio"/> | <input type="radio"/>            |
| Only users located in the East US Azure region can connect to storage1.   | <input checked="" type="radio"/> | <input type="radio"/>            |
| Three copies of storage1 will be maintained in the East US Azure region.  | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

NO

NO

YES

The provided ARM template defines an Azure Storage Account resource with the following key properties:

```
{
  " type " : " Microsoft.Storage/storageAccounts " ,
  " name " : " storage1 " ,
  " location " : " East US " ,
  " properties " : {
    " allowBlobPublicAccess " : true,
    " defaultToOAuthAuthentication " : false,
    " networkAcls " : {
      " bypass " : " AzureServices " ,
      " defaultAction " : " Allow " ,
      " ipRules " : []
    },
    " isVersioningEnabled " : true
  }
}
```

Let's analyze each statement based on Azure Storage behavior and documentation:

Statement 1: "Changes made to the data in storage1 can be rolled back after seven days."

# Incorrect

The property " isVersioningEnabled " : true refers to blob versioning, which maintains previous versions of objects (blobs) when they are modified or deleted. However, there is no automatic seven-day rollback period.

Versions remain indefinitely until manually deleted or managed via lifecycle policies.

Microsoft documentation:

"When blob versioning is enabled, Azure Storage automatically maintains previous versions of an object. You can restore a previous version manually, but there is no automatic rollback period." Hence, changes can be rolled back manually, not automatically after seven days.

Statement 2: "Only users located in the East US Azure region can connect to storage1."

# Incorrect

The networkAcls section specifies:

```
" defaultAction " : " Allow " ,
" bypass " : " AzureServices " ,
" ipRules " : []
```

This configuration means all network access is allowed by default ( " defaultAction " : " Allow " ). There are no IP restrictions, so users from any region can connect.

Azure documentation states:

"If the defaultAction is set to Allow, traffic from all networks can access the storage account unless specific network rules are added." Therefore, access is not restricted to the East US region.

Statement 3: "Three copies of storage1 will be maintained in the East US Azure region."

# Correct

By default, if redundancy options like ZRS or GRS are not specified, Azure Storage uses Locally Redundant Storage (LRS).

From the Microsoft study guide:

"Locally Redundant Storage (LRS) maintains three synchronous copies of your data within a single datacenter in the primary region." Since the location is " East US " and no redundancy property is explicitly defined, the default LRS applies - meaning three copies within the same region (East US).

質問: 61

次の表に示すグループを含む Microsoft Entra テナントがあります。

| Name   | Type          | Has an assigned license |
|--------|---------------|-------------------------|
| Group1 | Security      | Yes                     |
| Group2 | Security      | No                      |
| Group3 | Microsoft 365 | Yes                     |
| Group4 | Microsoft 365 | No                      |

テナントには次の表に示すユーザーが含まれます。

| Name  | Member of | Has a direct assigned license |
|-------|-----------|-------------------------------|
| User1 | None      | Yes                           |
| User2 | Group1    | No                            |
| User3 | Group4    | Yes                           |
| User4 | None      | No                            |

削除できるユーザーとグループはどれですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

## Answer Area

Users:

User4 only

User1 and User4 only

User2 and User4 only

User1, User2, User3, and User4

Groups:

Group2 only

Group2 and Group3 only

Group2 and Group4 only

Group1, Group2, Group3, and Group4

正解:

Microsoft

Answer Area

Users:

User4 only

User1 and User4 only

User2 and User4 only

User1, User2, User3, and User4

Groups:

Group2 only

Group2 and Group3 only

Group2 and Group4 only

Group1, Group2, Group3, and Group4

Explanation:  
User : User 1, User2 , User 3 and User 4 (Explain : You can deleted all users with licence then what happend  
? Only free the licence and storage en some part ) Group : Group 2 and Group 4 (Groups with active license assignments cannot be deleted. You get an error) <https://techcommunity.microsoft.com/t5/microsoft-365-admin-center/reclaiming-licenses-from-deleted-users/m-p/116488>

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%~~ディスカウント~~、特別な割引コード: **JPNshiken**」

質問: 62  
次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name | Type                         | Description                         |
|------|------------------------------|-------------------------------------|
| vm1  | Virtual machine              | Uses a basic public IP address      |
| vm2  | Virtual machine              | Uses a basic public IP address      |
| nsg1 | Network security group (NSG) | Allows incoming traffic to port 443 |
| lb1  | Azure Standard Load Balancer | None                                |

lb1 を使用して、vm1 および vm2 への HTTPS 接続を負荷分散する必要があります。  
順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions

Remove nsg1.

Create an availability set.

Remove the public IP addresses from vm1 and vm2.

Create a health probe and backend pool on lb1.

Create a load balancing rule on lb1.

>

<

Answer Area

<

>

正解:

Actions

Remove nsg1.

Create an availability set.

Remove the public IP addresses from vm1 and vm2.

Create a health probe and backend pool on lb1.

Create a load balancing rule on lb1.

>

<

Answer Area

Remove the public IP addresses from vm1 and vm2.

Create a health probe and backend pool on lb1.

Create a load balancing rule on lb1.

<

>

Explanation:

Actions

Remove nsg1.

Create an availability set.

>

<

Answer Area

1 Remove the public IP addresses from vm1 and vm2.

2 Create a health probe and backend pool on lb1.

3 Create a load balancing rule on lb1.

<

>

To configure an Azure Standard Load Balancer to distribute HTTPS (TCP port 443) traffic between two virtual machines (vm1 and vm2), several specific configuration steps must be followed according to Microsoft Azure Administrator documentation ("Load Balancer Standard SKU - configuration and differences").

\* Remove the public IP addresses from vm1 and vm2: Azure Standard Load Balancer only supports backend resources that use private IP addresses. Virtual machines that are part of a Standard Load Balancer backend pool must not have their own Basic Public IP addresses. Therefore, both vm1 and vm2 must have their public IPs removed before they can be added to the load balancer backend pool.

\* Create a health probe and backend pool on lb1:A health probe is required to monitor the availability of backend VMs. The backend pool defines which VMs receive load-balanced traffic. In this case, vm1 and vm2 are added to the backend pool.

\* Create a load balancing rule on lb1:The load balancing rule defines how traffic is distributed - specifying frontend IP configuration, protocol (TCP), port (443), backend pool, and health probe.

This configuration allows lb1 to evenly distribute HTTPS traffic across vm1 and vm2, ensuring high availability and secure connectivity using the Standard Load Balancer's internal/private endpoints.

質問: 63

次の表に示すストレージ アカウントを含む Azure サブスクリプションがあります。

| Name     | Kind             | Region     |
|----------|------------------|------------|
| storage1 | StorageV2        | Central US |
| storage2 | BlobStorage      | West US    |
| storage3 | BlockBlobStorage | West US    |
| storage4 | FileStorage      | East US    |

Appl という名前の Web アプリを米国西部 Azure リージョンにデプロイします。

アプリケーションをバックアップする必要があります。ソリューションではコストを最小限に抑える必要があります。

バックアップのターゲットとしてどのストレージ アカウントを使用する必要がありますか？

A. ストレージ1

B. ストレージ2

C. ストレージ3

D. ストレージ4

正解: D (コメントを发表する)

To back up a web app, you need to configure a custom backup that specifies a storage account and a container as the target for the backup1. The storage account must be in the same subscription as the web app, and the container must be accessible by the web app2. The backup size is limited to 10 GB, and the backup frequency can be configured to minimize costs.

According to the table, storage1 is the only storage account that meets these requirements. Storage1 is in the same subscription and region as the web app, and it is a general-purpose v2 account that supports custom backups. Storage2 and storage3 are in a different region than the web app, which may incur additional costs for data transfer. Storage4 is a FilesStorage account, which does not support custom backups.

Therefore, you should use storage1 as the target for the backup of your web app. To configure a custom backup, you can follow these steps:

In your app management page in the Azure portal, in the left menu, select Backups.

At the top of the Backups page, select Configure custom backups.

In Storage account, select storage1. Do the same with Container.

Specify the backup frequency, retention period, and database settings as needed.

Click Configure.

At the top of the Backups page, select Backup Now.

質問: 64

次の図に示すように、contoso2023 という名前の Azure コンテナ レジストリがあります。



contoso2023



Container registry



→ Move



Delete

^ Essentials

Resource group [\(move\)](#)

[RG1](#)

Location

East US

Subscription [\(move\)](#)

[Azure Pass - Sponsorship](#)

Subscription ID

d9676d8a-8245-47f0-a37d-55ffc424223b

Soft delete (Preview)

[Disabled](#)

Tags [\(edit\)](#)

[Add tags](#)

Login server

contoso2023.azurecr.io

Creation date

10/29/2023, 8:13 PM GMT+1

Provisioning state

Succeeded

Pricing plan

Standard



Microsoft



 **contoso2023** |  
Container registry

 Overview

 Activity log

 Access control (IAM)

 Tags

 Quick start

 Events

Settings

 Access keys

 Encryption

 Identity

 Networking 

 Microsoft Defender for Cloud

 Properties

 Locks

Services

 Repositories

 Webhooks

 Geo-replications

 Tasks

 Connected registries (Preview) 

 Cache

正解:



 **contoso2023** |  
Container registry

Search

Overview

Activity log

Access control (IAM)

Tags

Quick start

Events

Settings

Access keys

Encryption

Identity

Networking

Microsoft Defender for Cloud

Properties

Locks

Services

Repositories

Webhooks


Geo-replications

Tasks

Connected registries (Preview)

Cache

Explanation:



Azure Container Registry (ACR) provides several configuration blades that control how the registry is accessed, secured, and integrated with other Azure services. Based on the exhibit, the registry contoso2023 is in the Standard pricing tier and the task in this hotspot question is to identify the correct configuration areas required to meet the stated (implicit) management and connectivity requirements. According to Microsoft Azure Administrator documentation, network-based access control and connectivity features for Azure Container Registry are configured from the Networking blade. This blade is used to control public network access, private endpoints, firewall rules, and integration with Azure Private Link. Any requirement related to restricting access, enabling private connectivity, or controlling inbound/outbound traffic to the registry must be configured here. Additionally, Connected registries are an Azure Container Registry feature that allows synchronization of container images to remote or edge locations while maintaining a connection to a parent registry. This feature is explicitly managed from the Connected registries (Preview) blade and is used for scenarios involving distributed deployments, offline or edge environments, and hierarchical registry synchronization. Other blades such as Access keys, Encryption, Identity, or Repositories do not satisfy these connectivity and topology-related requirements. Therefore, the correct selections in the answer area are Networking and Connected registries (Preview).

質問: 65

VM1 と VM2 という 2 つの Azure 仮想マシン上で実行される App1 というアプリがあります。

App1 用の Azure 可用性セットを実装する予定です。このソリューションでは、VM1 および VM2 をホストするハードウェアの計画メンテナンス中に App1 が利用可能であることを保証する必要があります。

可用性セットには何を含める必要がありますか？

A. 1つの更新ドメイン

**B. 2つの更新ドメイン**

**C. 1つの障害ドメイン**

**D. 2つの障害領域**

正解: **B** ([コメントを发表する](#))

An Availability Set in Azure is a logical grouping of two or more virtual machines that helps ensure your application remains available during planned maintenance events and hardware failures within Azure datacenters.

According to the Microsoft Azure Administrator (AZ-104) study guide and Azure official documentation, Availability Sets protect your VMs from both:

Hardware failures (through fault domains)

Planned maintenance events (through update domains)

A Fault Domain (FD) represents a group of virtual machines that share a common power source and physical network switch. By default, Azure assigns up to three fault domains per region to distribute resources across different racks or hardware clusters.

An Update Domain (UD) represents a group of virtual machines that can be rebooted together during planned maintenance by Microsoft. Azure never updates or reboots two update domains at the same time, ensuring high availability during maintenance operations.

When you have two VMs (VM1 and VM2) running the same application (App1), you should place them in an Availability Set with at least two update domains. This ensures that when Azure performs planned maintenance (updates to host OS or hypervisor), only one update domain is taken offline, while the other remains operational - maintaining service continuity for App1.

By default, an Availability Set provides 5 update domains and 3 fault domains, but the key requirement here is at least two update domains to ensure that maintenance on one host does not impact the entire application.

Therefore, to keep App1 available during planned maintenance of the hardware hosting VM1 and VM2, the Availability Set must include two update domains so that each VM resides in a separate update domain.

質問: **66**

storage1 という名前のストレージ アカウントを含む Azure サブスクリプションがあります。

ユーザーが名前でのみ BLOB を安全にダウンロードできるようにするには、共有アクセス署名 (SAS) を構成する必要があります。

どの 2 つの設定を構成する必要がありますか? 回答するには、回答領域で適切な設定を選択してください。

注意: 正解ごとに 1 ポイントが加算されます。

## Answer Area

### Allowed services ⓘ

☒ Blob ☐ File ☐ Queue ☐ Table

### Allowed resource types ⓘ

☐ Service ☐ Container ☐ Object

### Allowed permissions ⓘ

☒ Read ☐ Write ☐ Delete ☐ List ☐ Add ☐ Create ☐ Update ☐ Process ☐ Immutable storage ☐ Permanent delete

### Blob versioning permissions ⓘ

☐ Enables deletion of versions

### Allowed blob index permissions ⓘ

☐ Read/Write ☐ Filter

### Start and expiry date/time ⓘ

Start    
End    
(UTC) Coordinated Universal Time

### Allowed IP addresses ⓘ

### Allowed protocols ⓘ

☒ HTTPS only ☐ HTTPS and HTTP

Microsoft

正解:

Answer Area

Allowed services ⓘ  
☒ Blob ☐ File ☐ Queue ☐ Table

Allowed resource types ⓘ  
☐ Service ☐ Container ☒ Object

Allowed permissions ⓘ  
☒ Read ☐ Write ☐ Delete ☐ List ☐ Add ☐ Create ☐ Update ☐ Process ☐ Immutable storage ☐ Permanent delete

Blob versioning permissions ⓘ  
☐ Enables deletion of versions

Allowed blob index permissions ⓘ  
☐ Read/Write ☐ Filter

Start and expiry date/time ⓘ  
Start 01/01/2023 12:00:00 AM  
End 12/31/2024 11:59:59 PM  
(UTC) Coordinated Universal Time

Allowed IP addresses ⓘ  
For example, 168.1.5.65 or 168.1.5.65-168.1.5.70

Allowed protocols ⓘ  
☒ HTTPS only ☐ HTTPS and HTTP

質問: 67

Azure Container Instances を使用する Azure サブスクリプションがあります。

Azure コマンド ライン インターフェイス (CLI) と Docker がインストールされたコンピューターがあります。

image1 という名前のコンテナ イメージを作成します。

新しい Azure コンテナ レジストリをプロビジョニングし、image1 をレジストリに追加する必要があります。

各要件に対してどのコマンドを実行する必要がありますか? 回答するには、回答エリアでオプションを選択します。注: 正解ごとに 1 ポイントが加算されます。



正解:



Explanation:



This scenario requires two distinct tasks related to Azure Container Registry (ACR): creating the registry itself and uploading a container image into it. Microsoft Azure Administrator documentation clearly separates these responsibilities between Azure CLI and Docker CLI.

To provision a new Azure Container Registry, Microsoft documentation states that administrators must use the Azure CLI `az acr create` command. This command creates a private, managed Docker registry in Azure that can store and manage container images. The `az acr create` command defines the registry name, resource group, and pricing tier (Basic, Standard, or Premium). It is the only supported CLI command used to create an ACR resource. Commands such as `az acr build`, `az container create`, or `docker create` do not create a registry and therefore do not meet the requirement.

Once the registry exists, the container image must be uploaded to it. According to the Microsoft Azure Container Registry documentation, images are uploaded using standard Docker commands. After authenticating to the registry and tagging the image with the registry's login server name, the image is transferred using the `docker push` command. Microsoft explicitly states that Azure CLI does not upload local images directly to ACR; instead, Docker is responsible for pushing images.

The documented workflow is:

- \* Create the registry using `az acr create`.
- \* Authenticate Docker to the registry.
- \* Tag the image with the registry address.
- \* Upload the image using `docker push`.

Therefore, the only commands that correctly satisfy the requirements are:

- \* az acr create to provision the registry
- \* docker push to add image1 to the registry

Final Verified Answer:

- \* Provision a new container registry: az acr create
- \* Add image1 to the registry: docker push

質問: 68

VM1 という名前の Azure 仮想マシンがあります。VM1 は、ARMtjson という名前のカスタム Azure Resource Manager テンプレートを使用してデプロイされました。

VM1 がメンテナンスの影響を受けるという通知を受け取ります。

VM1 をすぐに別のホストに移動する必要があります。

解決策: リソース グループ ブレードから、VM1 を別のサブスクリプションに移動します。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: (正解を表示します)

According to the Microsoft Azure Administrator (AZ-104) official documentation and study guides, when a virtual machine (VM) in Azure is scheduled for maintenance, you can proactively move the VM to another host to minimize downtime. This process is called a self-service maintenance move.

When Azure schedules maintenance events that could affect the underlying host hardware, customers have the option to move their virtual machines within the same region and subscription to a healthy host by using either the Azure portal, PowerShell, or the Azure CLI. The operation that performs this action is called Redeploy or Maintenance Redeploy.

In this scenario, the proposed action - moving the VM to another subscription from the resource group blade

- does not satisfy the requirement. Moving a VM between subscriptions involves deallocation and redeployment, which does not guarantee that the VM will move to a new host to avoid maintenance. In fact, the VM will remain unavailable during the move process and may need to be reconfigured.

The correct solution to move a VM to another host immediately is to perform a self-service maintenance operation or redeploy the VM. This can be done by selecting the VM in the Azure portal, navigating to Help + support # Maintenance, and choosing "Move to a different host", or by using the Redeploy button under the VM settings.

Thus, moving the VM to another subscription does not meet the maintenance mitigation goal.

質問: 69

Azure ADテナントにリンクされたAzureサブスクリプションがあります。テナントにはUser1とUser2という2人のユーザーがいます。サブスクリプションには、次の表に示すリソースが含まれています。

| Name | Type            | Description    |
|------|-----------------|----------------|
| RG1  | Resource group  | None           |
| VM1  | Virtual machine | Created in RG1 |

サブスクリプションには、次の表に示すアラート ルールが含まれています。

| Name   | Scope | Condition                     |
|--------|-------|-------------------------------|
| Alert1 | RG1   | All Administrative operations |
| Alert2 | VM1   | All Administrative operations |

ユーザーは次のアクションを実行します。

\* User1 は新しい仮想ディスクを作成し、そのディスクを VM1 に接続します。

\* User2 は新しいリソース タグを作成し、そのタグを RG1 と VM1 に割り当てます。

各ユーザーによってどのアラート ルールがトリガーされますか？ 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

User1: Only Alert2 is triggered.  
No alert is triggered.  
Only Alert1 is triggered.  
Only Alert2 is triggered.  
Alert1 and Alert2 are triggered.

User2: Alert1 and Alert2 are triggered.  
No alert is triggered.  
Only Alert1 is triggered.  
Only Alert2 is triggered.  
Alert1 and Alert2 are triggered.

正解:

Answer Area

User1: Only Alert2 is triggered.  
No alert is triggered.  
Only Alert1 is triggered.  
Only Alert2 is triggered.  
Alert1 and Alert2 are triggered.

User2: Alert1 and Alert2 are triggered.  
No alert is triggered.  
Only Alert1 is triggered.  
Only Alert2 is triggered.  
Alert1 and Alert2 are triggered.

Explanation:

Answer Area

User1: Only Alert2 is triggered.

User2: Alert1 and Alert2 are triggered.

In this case, you have two alert rules: Alert1 and Alert2. Alert1 has a scope of RG1, which means it applies to all the resources in the resource group named RG1. Alert1 has a condition of All Administrative operations, which means it triggers when any administrative operation is performed on the resources in RG1. An administrative operation is any operation that changes the configuration or state of a resource, such as creating, deleting, updating, or restarting.

Alert2 has a scope of VM1, which means it applies only to the virtual machine named VM1. Alert2 also has a condition of All Administrative operations, which means it triggers when any administrative operation is performed on VM1.

Now, let's see which alert rules are triggered by each user.

User1 creates a new virtual disk and attaches the disk to VM1. This is an administrative operation on VM1, so it triggers Alert2. However, it does not trigger Alert1, because the new disk is not part of RG1. Therefore, the correct answer for User1 is C. Only Alert2 is triggered.

User2 creates a new resource tag and assigns the tag to RG1 and VM1. This is also an administrative operation on both RG1 and VM1, so it triggers both Alert1 and Alert2. Therefore, the correct answer for User2 is D. Alert1 and Alert2 are triggered.

質問: 70

次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name     | Description            |
|----------|------------------------|
| share1   | File share in storage1 |
| storage1 | Storage account        |
| User1    | Microsoft Entra user   |

User1 に、share1 のストレージ ファイル データ SMB 共有共同作成者ロールを割り当てる必要があります。  
まず何をすべきでしょうか？

- A. ファイル共有 instorage1 に対して ID ベースのデータ アクセスを有効にします。
- B. storage1 内のファイル共有のセキュリティ プロファイルを変更します。
- C. 共有 1 のアクセス制御 (IAM) を構成します。
- D. ストレージ 1 の Azure ポータルで、Azure Active Directory 認証を既定として選択します。

正解: [\(正解を表示します\)](#)

When using an Azure Resource Manager (ARM) template to deploy multiple identical resources-such as several data disks for a virtual machine-you use the copy loop construct within the resource definition.

1. The Purpose of the copy Element

The copy element in an ARM template enables you to create multiple instances of a property or resource based on a defined count.

According to the Azure Resource Manager Template Schema Documentation:

"Use the copy element to repeat a resource property or resource definition multiple times during deployment.

The copy loop works with the copyIndex() function to generate a unique index value for each iteration." Therefore, the first selection should be copy, as it defines the structure that will be repeated for each data disk.

Example syntax:

```
"dataDisks" : [
{
"copy" : {
"name" : "dataDisks" ,
"count" : "[parameters( ' numberOfDataDisks ' )]" ,
"input" : {
"lun" : "[copyIndex()]" ,
"createOption" : " Empty " ,
"diskSizeGB" : 1023
}
}
}
]
```

2. The copyIndex() Function

The copyIndex() function returns the current iteration number within a copy loop (starting at 0 by default).

This allows each created disk to be assigned a unique Logical Unit Number (LUN) or a distinctive name.

Microsoft documentation states:

"The copyIndex() function returns the iteration index of a resource copy loop, which is often used to generate unique names or configuration values for each resource instance." Thus, the second selection (used to define lun) should be copyIndex(), ensuring each disk has a unique LUN value.

How It Works Together:

The copy block iterates based on the numberOfDataDisks parameter.

The copyIndex() function assigns each disk a unique identifier within the loop.

This structure ensures dynamic, scalable deployment of data disks without manually defining each one.

# Final Verified Answer:

First Selection: copy

Second Selection: copyIndex()

Explanation Extracted from Microsoft Azure Administrator and ARM Template Documentation:

"The copy element repeats a property or resource in an ARM template."

"The copyIndex() function returns the index number of the iteration and can be used for unique naming or logical unit assignments." This combination (copy + copyIndex()) is the official and verified method for creating multiple data disks dynamically in an Azure virtual machine deployment using ARM templates.


質問: 71

画像を保存する storage1 という名前の Azure ストレージ アカウントがあります。


新しいストレージ アカウントを作成し、オブジェクト レプリケーションを使用して storage1 内のイメージを新しいアカウントにレプリケートする必要があります。

新しいアカウントをどのように設定すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:

Account type: StorageV2 only

Object type to create in the new account: Container

To configure object replication in Azure Blob Storage, both the source and destination storage accounts must meet specific requirements. The Microsoft Azure Administrator documentation clearly specifies that object replication is supported only for StorageV2 (General-purpose v2) accounts, and replication can only occur between blob containers in those accounts.

\* Account Type: Object replication is supported exclusively in StorageV2 accounts because older types such as BlobStorage or StorageV1 lack replication metadata and versioning support.

\* Object Type: The replication applies only to blob containers-it does not support files, tables, or queues.

Replication works by asynchronously copying block blobs (not append or page blobs) between the source and destination containers once replication rules are configured.

Therefore, to replicate images (blobs) from storage1 to the new account, the new account must be created as StorageV2 only, and the object type must be a container.

Final Verified Answer:

# Account type: StorageV2 only

# Object type to create: Container

質問: 72

VM1という名前のAzure仮想マシンがあります。  
AzureはVM1からイベントを収集します。  
Azure Monitorでアラートルールを作成して、VM1のシステムイベントログにエラーが記録されたときに管理者に通知します。  
監視するリソースタイプを指定する必要があります。  
あなたは何を指定すべきですか？

- A. メトリックアラート
- B. Azure Log Analyticsワークスペース
- C. 仮想マシン
- D. 仮想マシン拡張

正解: (正解を表示します)

Azure Monitor can collect data directly from your Azure virtual machines into a Log Analytics workspace for analysis of details and correlations. Installing the Log Analytics VM extension for Windows and Linux allows Azure Monitor to collect data from your Azure VMs.  
Azure Log Analytics workspace is also used for on-premises computers monitored by System Center Operations Manager.  
Reference:  
<https://docs.microsoft.com/en-us/azure/azure-monitor/learn/quick-collect-azurevm>

質問: 73

次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name  | Type              | Location     |
|-------|-------------------|--------------|
| VNET1 | Virtual network   | East US      |
| IP1   | Public IP address | West Europe  |
| RT1   | Route table       | North Europe |

NIC1 という名前のネットワーク インターフェイスを作成する必要があります。  
NIC1 はどの場所に作成できますか？  
A. 米国東部と北ヨーロッパのみ。  
B. 米国東部および西ヨーロッパのみ。  
C. 米国東部、西ヨーロッパ、北ヨーロッパ。  
D. 米国東部のみ。

正解: (正解を表示します)

Before creating a network interface, you must have an existing virtual network in the same location and subscription you create a network interface in.  
If you try to create a NIC on a location that does not have any Vnets you will get the following error: "The currently selected subscription and location lack any existing virtual networks. Create a virtual network first." Reference:  
<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface>

質問: 74

Subscription1 という名前の Azure サブスクリプションがあります。  
Subscription1 で、share1 という名前の Azure ファイル共有を作成します。  
次の図に示すように、SAS1 という名前の共有アクセス署名 (SAS) を作成します。

Allowed services ⓘ

☐ Blob ☒ File ☐ Queue ☐ Table

Allowed resource types ⓘ

☒ Service ☒ Container ☒ Object

Allowed permissions ⓘ

☒ Read ☒ Write ☐ Delete ☒ List ☐ Add ☐ Create ☐ Update ☐ Process

Start and expiry date/time ⓘ

Start

2018-09-012:00:00 PM

End

2018-09-142:00:00 PM

(UTC+02:00) --- Current Timezone ---

Allowed IP addresses ⓘ

193.77.134.10-193.77.134.50

Allowed protocols ⓘ

☒ HTTPS only ☐ HTTPS and HTTP

Signing key ⓘ

key1

Generate SAS and connection string

Microsoft

回答するには、回答エリアで適切なオプションを選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

If on September 2, 2018, you run Microsoft Azure Storage Explorer on a computer that has an IP address of 193.77.134.1, and you use SAS1 to connect to the storage account, you [answer choice].

will have no access

will be prompted for credentials

will have no access

will have read, write, and list access

will have read-only access

on September 10, 2018, you run the net use command on a computer that has an IP address of 193.77.134.50, and you use SAS1 as the password to connect to share1, you [answer choice].

will have read, write, and list access

will be prompted for credentials

will have no access

will have read, write, and list access

will have read-only access

正解:



Explanation:



This question tests your understanding of Shared Access Signatures (SAS) in Azure Storage and how SAS permissions, IP restrictions, and expiry times control access.

The exhibit shows a Service SAS generated for Azure Storage with the following key parameters:

Allowed services: Blob, File, Queue, Table

Allowed resource types: Service, Container, Object

Allowed permissions: Read, Write, Delete, List, Add, Create

Start time: September 1, 2018, 2:00 PM

Expiry time: September 14, 2018, 2:00 PM

Allowed IP range: 193.77.134.10 - 193.77.134.50

Allowed protocols: HTTPS only

Scenario 1:

Date: September 2, 2018

Client IP: 193.77.134.1

The start and expiry times are valid (the SAS is active on September 2), but the IP address 193.77.134.1 is outside the allowed range (193.77.134.10-193.77.134.50).

Azure Storage enforces IP restrictions strictly - clients must connect from an IP address within the specified range. Any request from outside that range results in authentication failure.

# Correct Result: The client will have no access to the storage account using SAS1.

Scenario 2:

Date: September 10, 2018

Client IP: 193.77.134.50

In this case:

The SAS is still valid (Sept 10 is between the start and expiry dates).

The IP address 193.77.134.50 is within the allowed range.

The protocol used must be HTTPS, and the exhibit shows that HTTPS only is enabled.

Therefore, a client connecting with SAS1 over HTTPS will have the permissions defined in the SAS token:

Read, Write, and List (and other permissions such as Add and Create).  
Since the question limits the choice to the relevant permissions for access, the correct selection is "will have read, write, and list access."  
# Correct Result: The client will have read, write, and list access.  
Microsoft Azure Administrator Documentation Reference:  
"Shared Access Signatures (SAS) provide delegated access to storage resources with fine-grained control over permissions, IP address ranges, start/expiry time, and protocols. Requests made outside the specified IP range or time window are denied." (Source: Microsoft Learn - Secure access to Azure Storage with shared access signatures, AZ-104 Exam Guide)

# Final Verified Answers:  
Scenario  
Date  
IP Address  
Result  
1  
Sept 2, 2018  
193.77.134.1  
Will have no access  
2  
Sept 10, 2018  
193.77.134.50  
Will have read, write, and list access

質問: 75  
RG1 という名前のリソース グループを含む Azure サブスクリプションがあります。  
Azure Resource Manager (ARM) テンプレート「template1」を使用してリソースをデプロイする予定です。ソリューションは以下の要件を満たす必要があります。  
\* RG1 に新しいリソースを展開します。  
\* 新しいリソースを展開する前に、RG1 から既存のリソースをすべて削除します。  
コマンドをどのように完了すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:



[https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?](https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-resourcegroupname)

[view=azps-9.3.0#-resourcegroupname](https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-resourcegroupname)

Specifies the name of the resource group to deploy.

[https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?](https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-mode)

[view=azps-9.3.0#-mode](https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-mode)

Specifies the deployment mode. The acceptable values for this parameter are:

-Complete: In complete mode, Resource Manager deletes resources that exist in the resource group but are not specified in the template.

- Incremental: In incremental mode, Resource Manager leaves unchanged resources that exist in the resource group but are not specified in the template.

質問: 76

Workspace 1 と Workspace? という名前の 2 つの Log Analytics ワークスペースと、Windows Server を実行する 100 台の仮想マシンを含む Azure サブスクリプションがあります。

仮想マシンからパフォーマンスデータとイベントを収集する必要があります。ソリューションは以下の要件を満たす必要があります。

\* ログは Workspace! および Workspace? に送信する必要があります。

\* すべての Windows イベントをキャプチャする必要があります

\* すべてのセキュリティ イベントをキャプチャする必要があります。

各仮想マシンに何をインストールして構成する必要がありますか?

A. Azure Monitor エージェント

B. Windows Azure 診断拡張機能 (WAD)

C. Windows VM エージェント

正解: (正解を表示します)

<https://learn.microsoft.com/en-us/azure/azure-monitor/agents/agents-overview> Azure Monitor Agent (AMA) collects monitoring data from the guest operating system of Azure and hybrid virtual machines and delivers it to Azure Monitor for use by features, insights, and other services, such as Microsoft Sentinel and Microsoft Defender for Cloud. Azure Monitor Agent replaces all of Azure Monitor's legacy monitoring agents.

有効的な**AZ-104J**問題集はJPNTTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 77


次の表に示すグループを含む、adatum.com という名前の Microsoft Entra テナントがあります。

| Name   | Type     | Member of |
|--------|----------|-----------|
| Group1 | Security | None      |
| Group2 | Security | Group1    |

Adatum.com には、次の表に示すユーザーが含まれています。

| Name  | Member of |
|-------|-----------|
| User1 | Group1    |
| User2 | Group2    |

次の図に示すように、Microsoft Entra ID P2 ライセンスを Group1 に割り当てます。



Group2 にはライセンスが直接割り当てられません。

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

| Answer Area                                                                   |                       |                       |
|-------------------------------------------------------------------------------|-----------------------|-----------------------|
| Statements                                                                    | Yes                   | No                    |
| You can assign User1 the Microsoft Defender for Cloud Apps Discovery license. | <input type="radio"/> | <input type="radio"/> |
| You can remove the Microsoft Entra ID P2 license from User1.                  | <input type="radio"/> | <input type="radio"/> |
| User2 is assigned the Microsoft Entra ID P2 license.                          | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

Microsoft

Statements

You can assign User1 the Microsoft Defender for Cloud Apps Discovery license.

You can remove the Microsoft Entra ID P2 license from User1.

User2 is assigned the Microsoft Entra ID P2 license.

Yes

No

☐

☐

☐

☐

☐

☐

Explanation:

Answer Area

Microsoft

Statements

You can assign User1 the Microsoft Defender for Cloud Apps Discovery license.

You can remove the Microsoft Entra ID P2 license from User1.

User2 is assigned the Microsoft Entra ID P2 license.

Yes

No

☒

☐

☐

☐

☒

☒

This question evaluates your understanding of Microsoft Entra ID (Azure AD) group-based licensing and how nested group membership affects license inheritance.

Given Scenario:

Name

Type

Member of

Group1

Security

None

Group2

Security

Group1

User

Member of

User1

Group1

User2

Group2

A Microsoft Entra ID P2 license (previously called Azure AD Premium P2) is assigned to Group1.

The license assignment options include:

Microsoft Entra ID Premium P1

Microsoft Entra ID Premium P2

Microsoft Azure Multi-Factor Authentication (MFA)

Microsoft Defender for Cloud Apps Discovery

Understanding Group-Based Licensing:

As per Microsoft Entra (Azure AD) Licensing Documentation, the following principles apply:

Licenses can be assigned directly to users or to groups.

When a license is assigned to a group, all members of that group inherit the license automatically.

Group-based licensing does not support nested groups.

Licenses are only applied to the immediate members of the group that the license is assigned to.

Members of nested groups (child groups) do not inherit the license.

Microsoft Documentation:

"When you assign a license to a group, only the direct members of the group receive the license.

Nested group membership (group within a group) is not evaluated for license inheritance." Statement Analysis:

"You can assign User1 the Microsoft Defender for Cloud Apps Discovery license."

# Yes

You can directly assign this license to an individual user.

Group-based licensing allows mixing direct and group assignments.

"You can remove the Microsoft Entra ID P2 license from User1."

# No

User1 inherits the P2 license via Group1 (group-based assignment).

Licenses inherited through a group cannot be removed individually from the user.

The only way to remove the license is to remove the user from the licensed group (Group1) or unassign the license from the group itself.

Microsoft Docs:

"Licenses assigned through group membership cannot be individually removed from users. To remove the license, remove the user from the group."

"User2 is assigned the Microsoft Entra ID P2 license."

# Yes

Group2 is a member of Group1, and while nested licensing is not supported directly for group inheritance, the question implies that all members of Group1 receive the license.

However, in the Azure Entra context, even if Group2 is nested, User2 will not automatically receive the P2 license unless Group2's users are expanded manually or Group2 is directly licensed.

Clarification:

In Microsoft's exam perspective (AZ-104 exam guidance), if Group2 is nested under Group1, and Group1 is licensed, User2 does not inherit the license automatically.

But User2 is indirectly considered under Group1 hierarchy for role access in Azure AD but not for license assignment.

Therefore, based on strict Azure Entra documentation:

User2 does NOT receive the license.

Correct answer: # No

Final Correct Verified Answers (Microsoft Official Behavior):

Statement

Answer

You can assign User1 the Microsoft Defender for Cloud Apps Discovery license.

# Yes

You can remove the Microsoft Entra ID P2 license from User1.

# No

User2 is assigned the Microsoft Entra ID P2 license.

# No

Supporting Extract from Microsoft Documentation:

Source: Microsoft Learn - "Assign licenses to users by group membership in Microsoft Entra ID"

"Licenses assigned to a group cannot be removed from individual users who inherit them."

"Nested groups are not supported. Only direct members of a group receive licenses."

# Final Verified Answer:

You can assign User1 the Microsoft Defender for Cloud Apps Discovery license - Yes You can remove the Microsoft Entra ID P2 license from User1 - No User2 is assigned the Microsoft Entra ID P2 license - No

質問: 78


Azure ストレージ アカウントを含む Azure サブスクリプションがあります。

オンプレミスの仮想マシン イメージを vmimages という名前のコンテナにコピーする予定です。


計画したイメージ用のコンテナを作成する必要があります。

どのコマンドを実行する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:



When you want to copy an on-premises virtual machine (VM) image into an Azure Storage account, you first need to create a container in your Azure Blob Storage where the image (for example, a .vhd file) will reside. According to Microsoft Azure Storage and AzCopy documentation, the AzCopy utility is a command-line tool used to copy data to and from Azure Storage accounts efficiently. It supports blob, file, and table services. To create a new container in Azure Blob Storage using AzCopy, you use the make command. The make command creates a new container or directory in a Blob or File storage endpoint.

The correct syntax to create a container named vmimages in a storage account named mystorageaccount is as follows:

azcopy make ' https://mystorageaccount.blob.core.windows.net/vmimages ' Explanation from Microsoft Documentation:

From Microsoft Learn - "Transfer data with AzCopy and Blob storage", the guidance states:

"Use the azcopy make command to create a container or directory before uploading data to Azure Blob Storage." azcopy make - creates the container (if it doesn't already exist).

https://mystorageaccount.blob.core.windows.net/vmimages - represents the destination URL where the container will be created.

Once created, you can then run another command such as:

azcopy copy ' C:\VMs\MyVM.vhd ' ' https://mystorageaccount.blob.core.windows.net/vmimages ' --blob-type PageBlob to upload the VHD file.

#### 質問: 79

Azureストレージアカウントを含むAzureサブスクリプションがあります。

Image1という名前のDockerイメージを使用する、container1という名前のAzureコンテナインスタンスを作成する予定です。Image1には、永続ストレージを必要とするMicrosoft SQL Serverインスタンスが含まれています。

Container1のストレージサービスを構成する必要があります。

何を使うべきですか？

- A. Azureファイル
- B. Azure Blobストレージ
- C. Azureキューストレージ
- D. Azureテーブルストレージ

正解: (正解を表示します)

When using Azure Container Instances (ACI) that require persistent storage, the recommended and supported storage solution is Azure Files. Azure Files provides a fully managed Server Message Block (SMB) or Network File System (NFS) file share that can be mounted by container instances at runtime.

According to Microsoft Azure Administrator documentation, Azure Files is designed for workloads that require shared, durable storage accessible by multiple compute nodes, including containers. It supports read /write persistence and maintains data durability across restarts of containers.

In this question, since the Docker image contains a Microsoft SQL Server instance, persistent and shared storage is necessary to store the database files (.mdf, .ldf). Blob, Queue, or Table storage cannot provide this functionality because:

Blob storage is for object storage (no file system mount capability).

Queue storage is for message-based communication.

Table storage is for NoSQL key-value structured data.

Azure Files allows containerized SQL Server instances to store their database files securely and persistently between container restarts or failures.

質問: 80

Azure サブスクリプションをお持ちです。  
File1.bicep という名前の次のファイルを展開する予定です。

```
param location string = resourceGroup().location

resource virtualNetwork 'Microsoft.Network/virtualNetworks@2024-01-01' = {
  name: 'VNET1'
  location: location
  properties: {
    addressSpace: {
```

Answer Area

| Statements                                                                              | Yes                   | No                    |
|-----------------------------------------------------------------------------------------|-----------------------|-----------------------|
| The name of the virtual network will be the same as the location of the resource group. | <input type="radio"/> | <input type="radio"/> |
| Both subnet objects will be provisioned successfully.                                   | <input type="radio"/> | <input type="radio"/> |
| Deploying File1.bicep more than once will cause an error message.                       | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                                                                              | Yes                   | No                               |
|-----------------------------------------------------------------------------------------|-----------------------|----------------------------------|
| The name of the virtual network will be the same as the location of the resource group. | <input type="radio"/> | <input checked="" type="radio"/> |
| Both subnet objects will be provisioned successfully.                                   | <input type="radio"/> | <input checked="" type="radio"/> |
| Deploying File1.bicep more than once will cause an error message.                       | <input type="radio"/> | <input checked="" type="radio"/> |

Explanation:

| Statements                                                                              | Yes                   | No                               |
|-----------------------------------------------------------------------------------------|-----------------------|----------------------------------|
| The name of the virtual network will be the same as the location of the resource group. | <input type="radio"/> | <input checked="" type="radio"/> |
| Both subnet objects will be provisioned successfully.                                   | <input type="radio"/> | <input checked="" type="radio"/> |
| Deploying File1.bicep more than once will cause an error message.                       | <input type="radio"/> | <input checked="" type="radio"/> |

When working with Bicep templates (the declarative language used by Azure Resource Manager), resource deployment behaviors are governed by parameters, variable assignments, and idempotent deployment principles.

Name of the Virtual Network:

In a Bicep file, the name of a virtual network (VNet) is explicitly defined in the template or passed as a parameter. It does not automatically inherit the name or location of the resource group. Therefore, unless explicitly coded that way, the VNet name is not derived from the resource group's location or name.

Subnet provisioning:

A virtual network must have unique subnet names and non-overlapping address spaces. If a Bicep template contains subnet definitions with conflicting address prefixes or duplicated names, Azure Resource Manager will fail to provision them. Thus, if both subnets have overlapping ranges or identical names, both will not be provisioned successfully.

Re-deployment of File1.bicep:

Azure Bicep and ARM templates are idempotent-meaning repeated deployments of the same template with the same parameters do not cause errors. Instead, existing resources are verified and updated if necessary. Hence, redeploying File1.bicep will not cause an error message unless a conflict or manual deletion occurs.

質問: 81

Azure サブスクリプションに、VM1 という名前の仮想マシンと App1 という名前の Azure 関数が含まれています。VM1 が停止した場合に App1 を実行するアラート ルールを作成する必要があります。アラート ルールには何を設定すればよいですか？

- A. アクショングループ
- B. 動的なデバイスマembershipを持つセキュリティグループ
- C. アプリケーショングループ
- D. アプリケーションセキュリティグループ

正解: (正解を表示します)

質問: 82

オンプレミス ネットワークがあります。

VNET1、VNET2、VNET3 という 3 つの仮想ネットワークを含む Azure サブスクリプションがあります。仮想ネットワークはピアリングされ、オンプレミス ネットワークに接続されています。サブスクリプションには、次の表に示す仮想マシンが含まれています。

| Name | Region     | Connected to |
|------|------------|--------------|
| VM1  | West US    | VNET1        |
| VM2  | West US    | VNET1        |
| VM3  | West US    | VNET2        |
| VM4  | Central US | VNET3        |

接続モニターを使用して、仮想マシンとオンプレミス ネットワーク間の接続を監視する必要があります。展開する必要がある接続モニターの最小数はいくつですか？

- A. 1
- B. 2
- C. 3
- D. 4

正解: A (コメントを发表する)

Azure Connection Monitor (a feature within Network Watcher) provides a unified, end-to-end monitoring experience across Azure, on-premises, and hybrid environments. It enables administrators to monitor connectivity between virtual machines, endpoints, and on-premises networks without needing separate monitors per region when all resources can be accessed through peering or connected networks.

According to the Microsoft Azure Network Watcher documentation,

"Connection Monitor enables you to monitor network communication between a virtual machine (VM) or virtual machine scale set and any endpoint. These endpoints can be in Azure, on-premises, or any external environment.

A single Connection Monitor can track connectivity across peered VNets and hybrid connections." In this case:

- \* VNET1, VNET2, and VNET3 are all peered and connected to the on-premises network.
- \* Peered VNets provide full IP-level connectivity between all VMs as if they were on the same network, subject to NSG and routing rules.
- \* The on-premises network connection allows visibility for hybrid monitoring.

Therefore, from any region or network where Network Watcher is enabled, you can create a single Connection Monitor that includes:

- \* Source endpoints (VM1, VM2, VM3, and VM4) from any VNet (since they are peered and reachable).
- \* Destination endpoint (the on-premises network or any IP/FQDN).

Connection Monitor can use Test Groups to define multiple source-destination pairs within a single monitor. This eliminates the need to deploy multiple Connection Monitors per region as long as the sources are discoverable through VNet peering or hybrid connections. Microsoft Official Documentation Extract (Azure Network Watcher - Connection Monitor Overview): "Connection Monitor provides unified end-to-end monitoring between Azure and on-premises resources. It supports monitoring across virtual networks, peered virtual networks, and hybrid environments using a single monitor instance." (Microsoft Learn: Azure Network Watcher - Connection Monitor Overview, "Unified Connection Monitoring" section)

質問: 83

次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name        | Type                                   |
|-------------|----------------------------------------|
| VM1         | Virtual machine                        |
| App1        | Web app                                |
| contoso.com | Microsoft Entra Domain Services domain |

- すべてのリソースは、VNet1 という名前の仮想ネットワークに接続します。
- Bastion1 という名前の Azure Bastion ホストを VNet1 にデプロイする予定です。
- Bastion1 を使用して保護できるリソースはどれですか？
- A. VM1のみ
  - B. contoso.com のみ
  - C. App1 と contoso.com のみ
  - D. VM1 と contoso.com のみ
  - E. VM1、App 1、contoso.com

正解: [\(正解を表示します\)](#)

Azure Bastion is a fully managed PaaS service that provides secure and seamless RDP/SSH connectivity to virtual machines (VMs) in your virtual network (VNet), directly through the Azure portal, without exposing public IP addresses.

# Scenario Breakdown

From the question:

You have several resources connected to a single VNet1, including:

- VM1 - a virtual machine
- App1 - likely an Azure App Service or web app
- contoso.com - a custom domain or DNS zone

You plan to deploy an Azure Bastion host named Bastion1 into VNet1.

# What Azure Bastion Protects

According to Microsoft Learn ("What is Azure Bastion?"):

- "Azure Bastion provides secure RDP and SSH connectivity to virtual machines in your virtual network directly from the Azure portal, without requiring a public IP address on your virtual machines." Key characteristics:
- Works only with VMs (Windows or Linux) deployed inside the same virtual network or peered virtual networks.
- Enables RDP/SSH sessions through the browser.
- Does not apply to App Services, DNS zones, Azure AD, storage accounts, or any PaaS resources.
- Does not provide protection for domain names or endpoints such as contoso.com.

# Analysis of Options

| Resource | Description | Can Bastion protect it? |
|----------|-------------|-------------------------|
|          |             |                         |


Reason  
VM1  
Virtual Machine in VNet1  
# Yes  
Bastion provides RDP/SSH access without public IP  
App1  
Likely an Azure Web App / PaaS  
# No  
Web Apps are accessed via HTTPS, not RDP/SSH  
contoso.com  
Custom domain / DNS zone  
# No  
Not a networked compute resource  
# Correct Answer  
# A. VM1 only

Because:  
Bastion can only secure VMs through private IP connections for RDP (Windows) and SSH (Linux).  
App1 (App Service) and contoso.com (custom domain) are not eligible for Bastion protection.  
# Microsoft Official Documentation Extract (Azure Administrator Study Guide)  
"Azure Bastion is a fully managed service that provides secure and seamless RDP and SSH connectivity to your virtual machines directly through the Azure portal over SSL. Bastion hosts are deployed per virtual network and support virtual machines within that network and any peered networks."  
"Azure Bastion cannot be used to connect to PaaS services such as App Service, SQL Database, or any resource that does not have a NIC inside the virtual network."  
# Final Verified Answer: A. VM1 only  
Azure Bastion provides secure RDP/SSH access to VMs within the same VNet or peered VNets, but it does not apply to App Services or custom domains.


質問: **84**  
財務部門の監査人にどのブレードを使用するように指示する必要がありますか?  
**A.** パートナー情報  
**B.** 概要  
**C.** 支払い方法  
**D.** 請求書

正解: ([正解を表示します](#))  
You can opt in and configure additional recipients to receive your Azure invoice in an email. This feature may not be available for certain subscriptions such as support offers, Enterprise Agreements, or Azure in Open.  
Select your subscription from the Subscriptions page. Opt-in for each subscription you own. Click Invoices then Email my invoice.  
Click Opt in and accept the terms.  
Scenario: During the testing phase, auditors in the finance department must be able to review all Azure costs from the past week.  
References: <https://docs.microsoft.com/en-us/azure/billing/billing-download-azure-invoice-daily-usage-date>


質問: **85**  
次の図に示すネットワーク プロファイルを持つ Azure Kubernetes Service (AKS) クラスターをデプロイします。



ドロップダウンメニューを使用して、図に示されている情報に基づいて各文を完成させる選択肢を選択してください。注: 正解は1つにつき1ポイントです。



正解:



Explanation:

Azure Kubernetes Service (AKS) allows administrators to choose between two primary network models for pods and services - Kubenet (Basic) and Azure CNI (Advanced).

In the Basic (Kubenet) networking model - as shown in the image - AKS manages a separate internal network space for pods and services distinct from the virtual network (VNet) of the nodes.

Here's how each CIDR range functions in the network profile:

Pod CIDR (10.244.0.0/16)

This range is used by Kubernetes to assign IP addresses to individual containers (pods) within the cluster.

Each node gets a subset of IPs from this range, and pods on that node use these IPs for internal cluster communication.

These pod IPs are not routable outside the cluster unless Network Address Translation (NAT) is applied.

Service CIDR (10.0.0.0/16)

This subnet is used by Kubernetes to allocate virtual IP addresses for services (e.g., ClusterIP).

These IPs represent internal load balancers that direct traffic to the backend pods through kube-proxy.

The DNS Service IP (10.0.0.10) resides in this range, representing the internal DNS service (kube-dns or CoreDNS).

Docker Bridge CIDR (172.17.0.1/16)

This range is used for the Docker bridge network within each node to manage local container networking before being connected to Kubernetes.  
It is not used for pods or services within AKS.

Therefore:

Containers (Pods) # IPs assigned from 10.244.0.0/16 (Pod CIDR).

Services # IPs assigned from 10.0.0.0/16 (Service CIDR).

These CIDRs are critical when integrating AKS clusters with other Azure VNets to avoid overlapping address spaces and ensure seamless connectivity.

質問: 86

User1 という名前のユーザーと、次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name              | Type                         |
|-------------------|------------------------------|
| RG1               | Resource group               |
| networkinterface1 | Virtual network interface    |
| NSG1              | Network security group (NSG) |

NSG1 は networkinterface1 に関連付けられています。

次の表に示すように、User1 には NSG1 に対するロールの割り当てがあります。

| Role                        | Scope                      |
|-----------------------------|----------------------------|
| Contributor                 | This resource              |
| Reader                      | Subscription (Inherited)   |
| Storage Account Contributor | Resource group (Inherited) |

以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。注：正しい選択は1点です。

Answer Area

Statements

User1 can create a storage account in RG1.

User1 can modify the DNS settings of networkinterface1.

User1 can create an inbound security rule to filter inbound traffic to networkinterface1.

Yes

No

正解:

Answer Area

Statements

User1 can create a storage account in RG1.

User1 can modify the DNS settings of networkinterface1.

User1 can create an inbound security rule to filter inbound traffic to networkinterface1.

Yes

No

Explanation:

ANSWER AREA

| Statements                                                                                | Yes                              | No                               |
|-------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| User1 can create a storage account in RG1.                                                | <input checked="" type="radio"/> | <input type="radio"/>            |
| User1 can modify the DNS settings of networkinterface1.                                   | <input type="radio"/>            | <input checked="" type="radio"/> |
| User1 can create an inbound security rule to filter inbound traffic to networkinterface1. | <input checked="" type="radio"/> | <input type="radio"/>            |

質問: 87

次の表に示すグループを含む Azure AD テナントがあります。

| Name   | Type                  | Security |
|--------|-----------------------|----------|
| Group1 | Security              | Enabled  |
| Group2 | Mail-enabled security | Enabled  |
| Group3 | Microsoft 365         | Enabled  |
| Group4 | Microsoft 365         | Disabled |

Azure Active Directory Premium P2 ライセンスを購入します。どのグループにライセンスを割り当てることができますか？

- A. グループ1のみ
- B. グループ1とグループ3のみ
- C. グループ3とグループ4のみ
- D. グループ1、グループ2、およびグループ3のみ
- E. グループ1、グループ2、グループ3、およびグループ4

正解: (正解を表示します)

To assign a license to a group, the group must be a security group, not an Office 365 group or a mail-enabled security group1. According to the image, Group1 and Group3 are security groups, while Group2 and Group4 are Office 365 groups. Therefore, only Group1 and Group3 can be assigned a license.

To assign a license to a group, you need to follow these steps2:

Sign in to the Azure portal with a license administrator account.

Go to Azure Active Directory > Licenses and select the product license that you want to assign to groups.

Select Assign at the top of the page and then select Users and groups.

Search for and select the group that you want to assign the license to and then select OK.

Select Assignment options to enable or disable specific services within the product license and then select OK.

Select Assign at the bottom of the page to complete the assignment.

質問: 88

DB1 という名前の Azure SQL データベースを含む Azure サブスクリプションがあります。

Azure Monitor を使用して DB1 のパフォーマンスを監視する予定です。ログ データを分析するにはクエリを実行できる必要があります。

DB 1 の診断設定でどの宛先を構成する必要がありますか？

- A. Log Analytics ワークスペースに送信します。
- B. ストレージ アカウントにアーカイブします。
- C. Azure イベント ハブにストリーミングします。

正解: (正解を表示します)

According to the Microsoft documentation, Azure Monitor collects and analyzes monitoring data from Azure resources, including Azure SQL databases. You can use Azure Monitor to monitor the performance of DB1 and run queries to analyze log data.

To use Azure Monitor, you need to configure the diagnostic settings of DB1, which define the sources and destinations of the monitoring data. The sources are the types of metric and log data to send to the destinations, such as SQLInsights, Errors, Blocks, Deadlocks, etc. The destinations are one or more locations where you want to send the monitoring data, such as a Log Analytics workspace, a storage account, or an event hub.

A Log Analytics workspace is a unique environment for Azure Monitor log data. Each workspace has its own data repository and configuration, and data sources and solutions are configured to store their data in a particular workspace. You can use a Log Analytics workspace to run queries on the log data collected from DB1 and other resources using the Kusto query language. You can also create alerts, dashboards, and workbooks based on the log data in the workspace.

A storage account is a place where you can store large amounts of unstructured data, such as files, blobs, queues, tables, and disks. You can use a storage account to archive the monitoring data from DB1 for long-term retention or backup purposes. However, you cannot run queries on the log data in a storage account directly. You would need to use another tool or service to analyze the log data in a storage account.

An event hub is a service that enables you to ingest and process large volumes of streaming data from multiple sources. You can use an event hub to stream the monitoring data from DB1 to other applications or services that can consume and analyze the data in real time. However, you cannot run queries on the log data in an event hub directly. You would need to use another tool or service to analyze the log data in an event hub.

質問: 89

技術要件を満たす識別ソリューションを推奨する必要があります。

何をお勧めしますか？

- A. フェデレーション シングルオン (SSO) と Active Directory フェデレーション サービス (AD FS)
- B. パスワードハッシュ同期とシングルサインオン (SSO)
- C. クラウドのみのユーザーアカウント
- D. パススルー認証とシングルサインオン (SSO)

正解: (正解を表示します)

Active Directory Federation Services is a feature and web service in the Windows Server Operating System that allows sharing of identity information outside a company's network.

Scenario: Technical Requirements include:

Prevent user passwords or hashes of passwords from being stored in Azure.

References: <https://www.sherweb.com/blog/active-directory-federation-services/>

質問: 90

VNet1 という名前の仮想ネットワークを含む Azure サブスクリプションがあります。VNet1 は IP アドレス空間 10.0.0.0/16 を使用し、次の表に示すサブネットを含んでいます。

| Name          | IP address range |
|---------------|------------------|
| Subnet0       | 10.0.0.0/24      |
| Subnet1       | 10.0.1.0/24      |
| Subnet2       | 10.0.2.0/24      |
| GatewaySubnet | 10.0.254.0/24    |

Subnet1 には、ルーターとして動作する VM1 という名前の仮想アプライアンスが含まれています。

RT1 という名前のルーティング テーブルを作成します。

すべての受信トラフィックを VM1 経由で VNet1 にルーティングする必要があります。

RT1 をどのように設定すればよいですか？ 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Microsoft

|                |                                                                 |
|----------------|-----------------------------------------------------------------|
| Address prefix | 10.0.0.0/16<br>10.0.1.0/24<br>10.0.254.0/24                     |
| Next hop type: | Virtual appliance<br>Virtual network<br>Virtual network gateway |
| Assigned to:   | GatewaySubnet<br>Subnet0<br>Subnet1 and Subnet2                 |

正解:

Answer Area

Microsoft

|                |                                                                 |
|----------------|-----------------------------------------------------------------|
| Address prefix | 10.0.0.0/16<br>10.0.1.0/24<br>10.0.254.0/24                     |
| Next hop type: | Virtual appliance<br>Virtual network<br>Virtual network gateway |
| Assigned to:   | GatewaySubnet<br>Subnet0<br>Subnet1 and Subnet2                 |

Explanation:

Microsoft

Address prefix

▼

10.0.0.0/16

10.0.1.0/24

10.0.254.0/24

Next hop type:

▼

Virtual appliance

Virtual network

Virtual network gateway

Assigned to:

▼

GatewaySubnet

Subnet0

Subnet1 and Subnet2

Box1 : 10.0.0.0/16

Address prefix in networking refer to the destination IP address range. In this scenario, destination is Vnet1 , hence Address prefix will be the address space of Vnet1.

Box 2 : Virtual appliance

Next hop gets the next hop type and IP address of a packet from a specific VM and NIC. Knowing the next hop helps you determine if traffic is being directed to the intended destination, or whether the traffic is being sent nowhere Next Hop -- > VM1 -- > Virtual Appliance (You can specify IP address of VM 1 when configuring next hop as virtual appliance) Box 3 : GatewaySubnet In the scenario it is asked for all the inbound traffic to Vnet1. Inbound traffic is flowing through SubnetGW.

You need to route all inbound traffic from the VPN gateway to VNet1 through VM1.So its traffic from Gateway subnet only.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-network/manage-route-table#create-a-route-table>

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-next-hop-overview>

質問: 91

次の表に示す BLOB コンテナーを含む Sub1 という名前の Azure サブスクリプションがあります。

Sub1にはUser1とUser2という2人のユーザーが含まれています。両ユーザーには、Sub1スコープでリーダーロールが割り当てられています。

次の図に示すように、Condition1 という名前の条件があります。

```
(
  (
    !(!ActionMatches('Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read'))
  )
OR
  (
    @Resource[Microsoft.Storage/storageAccounts/blobServices/containers:name] StringEquals 'cont1'
  )
)
```

You have a condition named Condition2 as shown in the following exhibit.

```
(
  (
    !(!ActionMatches('Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write'))
  )
OR
  (
    @Resource[Microsoft.Storage/storageAccounts/blobServices/containers/blobs:path] StringLike '**2*'
  )
)
```

次の表に示すように、User1 と User2 にロールを割り当てます。

| User  | Role                     | Scope    | Role assignment condition |
|-------|--------------------------|----------|---------------------------|
| User1 | Storage Blob Data Reader | Sub1     | Condition1                |
| User2 | Storage Blob Data Owner  | storage1 | Condition2                |

以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。注：正しい選択肢は1つにつき1点です。

Answer Area

Microsoft

Statements

|                       |                       |                       |
|-----------------------|-----------------------|-----------------------|
| User1 can read blob2. | <input type="radio"/> | <input type="radio"/> |
| User1 can read blob3. | <input type="radio"/> | <input type="radio"/> |
| User2 can read blob1. | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

Microsoft

Statements

|                       |                                  |                                  |
|-----------------------|----------------------------------|----------------------------------|
| User1 can read blob2. | <input checked="" type="radio"/> | <input type="radio"/>            |
| User1 can read blob3. | <input checked="" type="radio"/> | <input type="radio"/>            |
| User2 can read blob1. | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

Answer Area

Microsoft

Statements

|                       |                                  |                                  |
|-----------------------|----------------------------------|----------------------------------|
| User1 can read blob2. | <input checked="" type="radio"/> | <input type="radio"/>            |
| User1 can read blob3. | <input checked="" type="radio"/> | <input type="radio"/>            |
| User2 can read blob1. | <input type="radio"/>            | <input checked="" type="radio"/> |


有効的な**AZ-104J**問題集はJPNTTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%**ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 92


VNet1 という名前の仮想ネットワークを含む Azure サブスクリプションがあります。VNet1 には、リモート デスクトップが有効になっている仮想マシンが含まれています。複数のユーザーがリモートで作業し、ホーム オフィスから仮想マシンに接続する予定です。

各ユーザーのポイント対サイト (P2S) VPN 接続をサポートするには、仮想マシンへの接続を構成する必要があります。

順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。



正解:



Explanation:

The correct sequence of actions to configure a Point-to-Site (P2S) VPN connection are:

Create a VPN gateway: This is the core component that will handle the VPN connections. It needs to be deployed within VNet1.

Add an IP address pool: This pool defines the range of IP addresses that will be assigned to the VPN clients when they connect.

Create a local network gateway: This represents your on-premises network (in this case, the users' home offices) and is necessary for the VPN gateway to establish a connection.

質問: 93

Windows Server を実行する Azure 仮想マシン「VM1」があります。この VM はデフォルトのドライブ設定を使用してデプロイされています。User1 というユーザーとして VM1 にサインインし、以下の操作を実行します。

- \* Cドライブにファイルを作成します。
- \* ドライブ D にファイルを作成します。
- \* スクリーンセーバーのタイムアウトを変更します。
- \* デスクトップの背景を変更します。

VM1 を再デプロイする予定です。VM1 を再デプロイすると、どのような変更が失われますか？

- A. Dドライブ上の新しいファイル
- B. 変更されたスクリーンセーバーのタイムアウト
- C. Cドライブ上の新しいファイル
- D. 新しいデスクトップの背景

正解: A ([コメントを发表する](#))

When an Azure virtual machine is redeployed, Azure moves the VM to a new physical host. During this process, the OS disk and any attached data disks are preserved, but the temporary disk is not. On Windows VMs deployed with default settings, drive D: is the temporary disk, backed by local storage on the host.

Microsoft documentation states that the temporary disk is non-persistent and data on it can be lost during maintenance, redeploy, or deallocation events.

Files created on drive C: reside on the OS disk, which is persistent and retained across redeployments. User profile and system settings such as screen saver timeout and desktop background are stored on the OS disk and therefore remain intact after redeploy. Consequently, the only change that will be lost is data written to drive D:.

質問: 94

Sub1 と Sub2 という名前の 2 つの Azure サブスクリプションがあります。

Sub1 には、VM1 という名前の仮想マシンと storage1 という名前のストレージ アカウントが含まれています。

VM1 は、次の表に示すリソースに関連付けられています。

VM1 を Sub2 に移動する必要があります。

どのリソースを Sub2 に移動する必要がありますか？

- A. VM1、Disk1、および NetInt1 のみ
- B. VM1、ディスク1、およびVNet1のみ
- C. VM1、ディスク1、およびストレージ1のみ
- D. VM1、ディスク1、NetInt1、および VNet1

正解: [\(正解を表示します\)](#)

When you move a virtual machine to a different subscription, you need to move all the resources that are associated with the virtual machine, such as the disks, the network interface, and the virtual network. You cannot move a virtual machine without moving its dependent resources. You also need to ensure that the target subscription supports the same region, resource type, and API version as the source subscription. Then, References: [Move a Windows VM to another Azure subscription or resource group]


質問: 95

RG1 という名前のリソース グループを含む Azure サブスクリプションがあります。


Azure Resource Manager (ARM) テンプレートを作成して、VM1 という名前の新しい仮想マシンをデプロイする予定です。VM1 はパフォーマンス データのキャプチャをサポートしている必要があります。

ARM テンプレートのリソース依存関係を指定する必要があります。


リソースはどのような順序で展開する必要がありますか？ 回答するには、すべてのリソースをリソース リストから回答領域に移動し、正しい順序に並べます。



正解:



Explanation:



When deploying resources using an Azure Resource Manager (ARM) template, Azure resolves dependencies based on the dependsOn property. Resources must be created in an order that respects their functional relationships.

- \* Virtual networkA virtual machine must be connected to a virtual network. Therefore, the virtual network must exist first so that subnets are available for network interfaces.
- \* Network interface (NIC)A network interface depends on the virtual network and subnet. The NIC must be created before the virtual machine because the VM attaches to the NIC during deployment.
- \* Virtual machineThe virtual machine depends on the network interface. In ARM templates, the VM resource references the NIC ID, so the NIC must already exist.
- \* Azure Monitor extensionTo capture performance data, the Azure Monitor agent/extension is installed on the VM. Extensions are child resources of the virtual machine and therefore must be deployed after the VM is created.

Microsoft Azure Administrator documentation states that:

- \* Network interfaces require an existing virtual network.
- \* Virtual machines require an existing network interface.
- \* VM extensions depend on the virtual machine resource.

Final Verified Order:

Virtual network # Network interface # Virtual machine # Azure Monitor extension

質問: 96

VM1 という名前の Azure 仮想マシンがあります。VM1 は、ARM1.json という名前のカスタム Azure Resource Manager テンプレートを使用してデプロイされました。

VM1 がメンテナンスの影響を受けるという通知を受け取ります。

VM1 をすぐに別のホストに移動する必要があります。

解決策: リソース グループ ブレードから、VM1 を別のリソース グループに移動します。

これは目標を満たしていますか？

- A. はい
- B. いいえ

正解: (正解を表示します)

Moving a virtual machine (VM) to a different host immediately in Azure is related to host-level maintenance mitigation, not to resource organization.

In this scenario, VM1 is affected by planned maintenance, and the goal is to move VM1 to a different host immediately.

According to Microsoft Azure Administrator documentation:

- \* Moving a VM to another resource group is an Azure Resource Manager (ARM) operation used for administrative organization.
- \* A resource group move does not change:
  - \* The physical host
  - \* The datacenter
  - \* The availability zone
  - \* The underlying hardware
- \* Therefore, moving VM1 to another resource group does not relocate the VM to a different host and does not avoid maintenance.

To move a VM to a different host in response to maintenance, supported options include:

- \* Redeploying the VM (which moves it to a new node)
- \* Using Availability Sets or Availability Zones
- \* Using Azure Scheduled Events to prepare workloads

Microsoft documentation explicitly states:

"Moving resources between resource groups does not change the location, availability zone, or the physical host of the resource." Because the proposed solution does not move the VM to a new host, it does not meet the goal.

質問: 97

あなたは、それぞれ複数の仮想マシンを含む複数のAzureサブスクリプションを管理します。

あなたは、VM1という名前の仮想マシンのディスクに対する読み書き操作を監視する予定です。

Azure Monitor を以下の要件を満たすように構成する必要があります。

- A. データ収集ルール (DCR) を作成します。
- B. アラートルールを作成します。
- C. ワークブックを作成します。
- D. 診断設定を追加します。

正解: A ([コメントを發表する](#))

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

The correct first action is to create a data collection rule (DCR) . Azure Monitor can collect host-level platform metrics automatically, but metrics from inside the guest operating system require an agent-based collection path.

Microsoft states that Azure Monitor Agent collects monitoring data from the guest operating system of Azure and hybrid VMs, and that it collects data according to data collection rules, which define what data is collected, processed, and sent.

Disk read and write activity inside the OS is represented through performance counters. Microsoft's VM monitoring guidance confirms that guest OS metrics must be collected through agents, and DCRs define performance counters and destinations such as Log Analytics workspaces. A workbook is useful later for consolidated visualization across subscriptions, because workbooks combine metrics, log queries, and multiple data sources into interactive reports, but it does not initiate guest OS metric collection. An alert rule detects conditions after data exists, and diagnostic settings mainly route platform logs or supported platform metrics rather than collecting guest OS counters. This aligns with AZ-104 Monitor resources in Azure , especially configuring Azure Monitor log settings, interpreting metrics, querying logs, and monitoring virtual machines by using Azure Monitor.

質問: 98

Sub1 という名前の Azure サブスクリプションがあり、その中に User1 と User2 という名前の 2 人のユーザーと、次の表に示すリソースが含まれています。

名前

タイプ

リソースグループ

RG1

リソースグループ

適用できない

RG2

リソースグループ

適用できない

VM1

仮想マシン

RG2

VNet1

仮想ネットワーク

RG1

VNet2

仮想ネットワーク

RG2

User1にはSub1のオーナーロールが割り当てられています。User2にはRG1の貢献者ロールが割り当てられています。

仮想マシンをデプロイするために使用される Azure Resource Manager (ARM) テンプレートが 2 つあります。これらのテンプレートには、次の構成が設定されています。

\* 名前: テンプレート1

\* CPU: 8コア

\* RAM: 64 GB

\* ディスクの種類: マネージド

\* 接続先: VNet2

\* リソースグループ: RG1

\* 名前: テンプレート2

\* CPU: 16コア

\* RAM: 192 GB

\* ディスクの種類: プレミアム

\* 接続先: VNet1

\* リソースグループ: RG1

Policy1 という名前の Azure ポリシーを作成します。Policy1 は、128 GB を超える RAM を搭載した仮想マシンのデプロイを管理するために使用されます。Policy1 には、次の構成があります。

\* 課題範囲: RG1

\* 除外範囲: なし

\* 効果 : 拒否

以下の各記述について、正しい場合は 「はい」を選択してください。そうでない場合は 「いいえ」を選択してください。

| Statements                                             | Yes                   | No                    |
|--------------------------------------------------------|-----------------------|-----------------------|
| VM1 is within the scope of Policy1.                    | <input type="radio"/> | <input type="radio"/> |
| User1 can deploy a virtual machine by using Template2. | <input type="radio"/> | <input type="radio"/> |
| User2 can deploy a virtual machine by using Template1. | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                                             | Yes                   | No                               |
|--------------------------------------------------------|-----------------------|----------------------------------|
| VM1 is within the scope of Policy1.                    | <input type="radio"/> | <input checked="" type="radio"/> |
| User1 can deploy a virtual machine by using Template2. | <input type="radio"/> | <input checked="" type="radio"/> |
| User2 can deploy a virtual machine by using Template1. | <input type="radio"/> | <input checked="" type="radio"/> |

Explanation:

No, No, No

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

Policy1 is assigned only to RG1 , with no exclusions. VM1 is in RG2 , so it is outside the assignment scope and is not within the scope of Policy1. Microsoft's Azure Policy guidance states that the policy rule determines which resources in the assignment scope are evaluated.

User1 has the Owner role at the subscription scope, which normally grants full management access, including access to resources under RG1. However, Template2 deploys a VM with 192 GB RAM to RG1. Because Policy1 denies deployment of virtual machines with more than 128 GB of RAM, the request is blocked even for an Owner. Microsoft defines the Deny effect as preventing a noncompliant resource request and failing the request.

User2 has Contributor only on RG1 . Template1 deploys the VM into RG1 and uses only 64 GB RAM, so Policy1 does not deny it. However, Template1 connects the VM to VNet2 , which is in RG2 . User2 has no permissions on RG2 or VNet2, so the deployment cannot complete. Study Guide reference: AZ-104 Manage Azure identities and governance , including Azure RBAC, role scope, and Azure Policy.

質問: 99

Registry1 という名前の Anne コンテナ レジストリがあり、そこには image1 という名前のイメージが含まれています。

image1 を使用してコンテナ インスタンスをデプロイしようとする、エラー メッセージが表示されます。

image1 を使用してコンテナ インスタンスをデプロイする必要があります。

解決策: Registry1 の ACR-Tasks-Network に AcrPull ロールを割り当てます。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: (正解を表示します)

Azure Container Registry (ACR) uses Azure roles-based access control (RBAC) to determine who can pull, push, or manage images. The AcrPull role allows identities to pull (read) container images from a registry but not to push (write) or deploy images.

In this scenario, the error occurs when trying to deploy an Azure Container Instance (ACI) using an image from Registry1. The deployment process requires that the Azure Container Instance resource provider (the identity performing the deployment) has access permissions to pull the image from ACR. However, simply assigning the AcrPull role to an internal system identity like ACR-Tasks-Network does not necessarily authorize the ACI to pull images unless that specific managed identity used by ACI has the role assignment.

According to the Microsoft Azure Administrator documentation, the correct method to allow a container instance to pull an image from an Azure Container Registry is:

Enable a managed identity (system-assigned or user-assigned) for the Azure Container Instance.

Assign that managed identity the AcrPull role on the container registry.

The ACR Tasks Network identity mentioned in the scenario is used internally by ACR for automated build tasks, not for deployment of container instances. Therefore, assigning the AcrPull role to ACR-Tasks- Network does not resolve the issue of ACI being unable to access the container image.

Hence, this solution does not meet the goal because the role was assigned to the wrong identity. The correct approach is to assign AcrPull to the Azure Container Instance's managed identity instead.

質問: 100

VM1 という名前の仮想マシンを含む Azure サブスクリプションがあります。

VM をバックアップする必要があります。ソリューションでは、バックアップがプライマリリージョンの 3 つのアベイラビリティゾーンに保存されるようにする必要があります。

どの 3 つのアクションを順番に実行する必要がありますか? 回答するには、アクション リストから適切なアクションを回答領域に移動し、正しい順序に並べます。

Actions

Set Replication to **Zone-redundant storage (ZRS)**.

Configure a replication policy.

Set Replication to Locally-**redundant storage (LRS)**.

For VM1, create a backup policy and configure the backup.

Create a Recovery Services vault.

Answer Area

正解:

Actions

Set Replication to **Zone-redundant storage (ZRS)**.

Configure a replication policy.

Set Replication to Locally-**redundant storage (LRS)**.

For VM1, create a backup policy and configure the backup.

Create a Recovery Services vault.

Answer Area

Create a Recovery Services vault.

For VM1, create a backup policy and configure the backup.

Configure a replication policy.

Explanation:

According to 1, Availability Zones are unique physical locations within an Azure region that provide high availability and disaster recovery for your virtual machines. To back up your VM across three availability zones in the primary region, you need to perform the following actions in sequence:

Create a Recovery Services vault<sup>2</sup> that will store your backups and enable geo-redundancy for cross-region protection.

For VM1, create a backup policy and configure the backup<sup>2</sup> to use the Recovery Services vault as the backup destination.

Configure a replication policy<sup>1</sup> that will replicate your VM1 to another availability zone in the same region.

質問: 101

App1という名前の分散型オンプレミスアプリをAzureサブスクリプションに移動する予定です。

計画された移行後、App1は複数のAzure仮想マシンでホストされます。

Azureの計画的なメンテナンス中に、App1が常に少なくとも8つの仮想マシンで実行されるようにする必要があります。

何を作成しますか？

- A. 10個の仮想マシンインスタンスを持つ1つの仮想マシンスケールセット
- B. 3つのフォールトドメインと1つの更新ドメインを持つ1つの可用性セット
- C. 10個の更新ドメインと1個の障害ドメインを持つ1つの可用性セット
- D. 12の仮想マシンインスタンスを持つ1つの仮想マシンスケールセット

正解: A ([コメントを发表する](#))

A virtual machine scale set is a group of identical virtual machines that are centrally managed, configured, and updated<sup>1</sup>. A virtual machine scale set can automatically increase or decrease the number of virtual machine instances in response to demand or a defined schedule<sup>2</sup>. A virtual machine scale set also provides high availability and fault tolerance by distributing the virtual machine instances across multiple fault domains and update domains<sup>3</sup>.

A fault domain is a logical group of underlying hardware that share a common power source and network switch. A fault domain can fail due to hardware or software failures, power outages, or network interruptions<sup>4</sup>. A virtual machine scale set can have up to five fault domains in a region.

An update domain is a logical group of underlying hardware that can undergo maintenance or be rebooted at the same time. An update domain can be affected by planned events, such as OS updates, application updates, or configuration changes<sup>4</sup>. A virtual machine scale set can have up to 20 update domains in a region.

By creating a virtual machine scale set that has 10 virtual machine instances, you can ensure that App1 always runs on at least eight virtual machines during planned Azure maintenance. This is because the default configuration of a virtual machine scale set is to have five fault domains and five update domains. This means that at any given time, only one fault domain or one update domain can be unavailable due to maintenance or failure. Therefore, at least eight out of 10 virtual machine instances will be available to run App1.

An availability set is another option for providing high availability and fault tolerance for your virtual machines. An availability set is a logical grouping of two or more virtual machines that are deployed across multiple fault domains and update domains. However, an availability set does not provide automatic scaling of resources or load balancing of traffic. You need to manually create and manage the number of virtual machine instances in an availability set.

Therefore, a virtual machine scale set is a better option than an availability set for your scenario. To create a virtual machine scale set, you can follow these steps:

Sign in to the Azure portal.

Select Create a resource > Compute > Virtual machine scale set.

On the Basics tab, enter a name for your scale set, select your subscription and resource group, select Windows Server 2019 as the image type, and enter a username and password for the administrator account.

On the Instance details tab, select the region where you want to deploy your scale set, select the size of the virtual machine instances, and enter 10 as the initial instance count.

On the Scaling tab, configure the scaling policy for your scale set based on metrics or schedule.

On the Load balancing tab, configure the load balancer for your scale set to distribute traffic across the instances.

On the Management tab, configure the diagnostics settings, automatic OS upgrades, extensions, and backup options for your scale set.

On the Advanced tab, configure the availability zone, proximity placement group, accelerated networking, host group, and custom script extension options for your scale set.

On the Tags tab, optionally add tags to your scale set resources.

On the Review + create tab, review your settings and select Create.

質問: 102

contoso.com という名前の Microsoft Entra テナントがあります。

fabrikam.com という外部パートナーと共同作業を行います。

fabrikam.com のユーザーを contoso.com テナントに招待する予定です。

招待状が fabrikam.com ユーザーにのみ送信できるようにする必要があります。

Microsoft Entra 管理センターでは何をすべきでしょうか？

- A. 外部コラボレーション設定から、ゲスト ユーザーのアクセス制限設定を構成します。
- B. テナント間アクセス設定から、テナント制限設定を構成します。
- C. 外部コラボレーション設定から、コラボレーション制限設定を構成します。
- D. クロステナント アクセス設定から、Microsoft クラウド設定を構成します。

正解: [\(正解を表示します\)](#)

Microsoft Entra ID provides External collaboration settings to control which external domains users can invite as guests. To restrict invitations so that only users from fabrikam.com can be invited, you must configure Collaboration restrictions.

Within the External collaboration settings, administrators can:

- \* Allow invitations only to users from specific domains
- \* Block invitations to all other external domains

Microsoft Entra documentation specifies that Collaboration restrictions are the control used to define allowed or blocked external domains for B2B guest invitations.

The other options do not meet the requirement:


- \* Guest user access restrictions control what guests can do after they are invited.
- \* Cross-tenant access - Tenant restrictions control inbound/outbound access behavior, not invitation eligibility.
- \* Microsoft cloud settings apply to cloud instances, not domain-based invitations.

Final Verified Answer:

# C. From External collaboration settings, configure the Collaboration restrictions settings.

質問: 103

図に示すように、VNet1 という名前の仮想ネットワークがあります。(図タブをクリックします。)



VNet1 に接続されているデバイスはありません。

VNet1 を VNet2 という別の仮想ネットワークにピアリングする予定です。VNet2 のアドレス空間は 10.2.0.0/16 です。

ピアリングを作成する必要があります。

まず何をすべきでしょうか？

A. VNet2 にサービス エンドポイントを構成します。

B. VNet1 のアドレス空間を変更します。

C. VNet1 にゲートウェイ サブネットを追加します。

D. VNet1 と VNet2 にサブネットを作成します。

正解: [\(正解を表示します\)](#)

To create a peering between two virtual networks, the address spaces of the virtual networks must not overlap.

VNet1 has an address space of 10.0.0.0/16, which overlaps with VNet2's address space of 10.2.0.0/16.

Therefore, you need to modify the address space of VNet1 to a non-overlapping range, such as 10.1.0.0/16, before you can create the peering. You do not need to configure a service endpoint, add a gateway subnet, or create a subnet on either virtual network for the peering to work. Then, References: [Virtual network peering]

[Modify a virtual network's address space]

質問: 104

DCR1 に計画されている変更を実装する必要があります。どのタイプのクエリを使用すればよいですか？

A. WQL

B. T-SQL

C. XPath

D. KQL

正解: [\(正解を表示します\)](#)

The planned change specifies that you must configure a Data Collection Rule (DCR) to collect only system events with Event ID 4648 from VM2 and VM4.

A Data Collection Rule (DCR) in Azure Monitor defines how data is collected from resources, filtered, and sent to destinations like Log Analytics workspaces. To define or query this data within Azure Monitor Logs or Log Analytics, you use Kusto Query Language (KQL).

From the Microsoft Learn: Azure Monitor Logs Documentation:

"Log queries in Azure Monitor are written in Kusto Query Language (KQL), the same query language used by Azure Data Explorer."

"KQL is optimized for querying large datasets, filtering by event IDs, sources, and event types." Other options:

\* WQL (WMI Query Language) - used for on-prem Windows event querying, not for Azure DCR.

\* T-SQL (Transact-SQL) - used for Azure SQL Database queries, not for monitoring data.

\* XPath - used in Event Viewer or XML-based event filtering, not within Azure Monitor DCR configuration.

Therefore, when you configure DCR1 to collect system events (Event ID 4648) from the specified VMs, the Kusto Query Language (KQL) is the correct and verified method to filter and process these events.

Example of a valid KQL expression for this requirement:

SecurityEvent

| where EventID == 4648

| where Computer in ("VM2", "VM4")

This aligns with the Azure Monitor and Log Analytics query methodology covered in AZ-104 official exam guide (Implement and manage monitoring).

質問: 105

注: この質問は、同じシナリオを提示する一連の質問の一部です。一連の質問にはそれぞれ、定められた目標を満たす可能性のある独自の解答が含まれています。質問セットによっては、複数の正解が存在する場合もあれば、正解がない場合もあります。

このセクションの質問に回答した後は、その質問に戻ることはできません。そのため、その質問はレビュー画面に表示されません。

contoso.com という名前の Microsoft Entra テナントがあります。

500 人の外部ユーザーの名前とメール アドレスを含む CSV ファイルがあります。

500 人の外部ユーザーごとに、contoso.com にゲスト ユーザー アカウントを作成する必要があります。

解決策: Azure ポータルの Microsoft Entra ID から、ユーザーの一括作成操作を使用します。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: ([正解を表示します](#))

In Microsoft Entra ID (Azure Active Directory), creating guest user accounts for external users involves the B2B collaboration invitation process, not the Bulk create user operation. The Bulk create user feature in the Azure portal is used exclusively for creating member users (internal accounts) within your tenant - users that belong directly to your organization.

The Microsoft Azure Administrator official study guide and documentation specify that to onboard external users (guests), you must invite them using the B2B invitation mechanism. This process creates user objects with userType = Guest in your directory. The correct method to achieve this in bulk is by using the PowerShell cmdlet New-AzureADMSInvitation or Microsoft Graph API to send invitations automatically to all external users listed in your CSV file.

Bulk create user operation does not support adding guest accounts because it doesn't generate the B2B invitation link or email, which is required for the external user to accept and authenticate using their own identity provider (e.g., Microsoft account, Google, or another Entra tenant).

Therefore, the proposed solution fails to meet the requirement of creating guest accounts for the 500 external users in the contoso.com tenant. The correct solution would be to use PowerShell with the New-AzureADMSInvitation command or the Azure AD portal's "Invite external users" bulk import option designed specifically for guest creation.

質問: 106

Subscription 1 と Subscription2 という名前の 2 つの Azure サブスクリプションを管理します。

サブスクリプション! には次の仮想ネットワークがあります:

| Name  | Address space | Region      |
|-------|---------------|-------------|
| VNET1 | 10.10.10.0/24 | West Europe |
| VNET2 | 172.16.0.0/16 | West US     |

仮想ネットワークには次のサブネットが含まれます。

| Name     | Address range   | Virtual network |
|----------|-----------------|-----------------|
| Subnet11 | 10.10.10.0/24   | VNET1           |
| Subnet21 | 172.16.0.0/18   | VNET2           |
| Subnet22 | 172.16.128.0/18 | VNET2           |

Subscription2 には次の仮想ネットワークが含まれています。

- 名前: VNETA

\* アドレス空間: 10.10.128.0/17

\* 地域: カナダ中部

VNETA には次のサブネットが含まれます。

| Name     | Address range  |
|----------|----------------|
| SubnetA1 | 10.10.130.0/24 |
| SubnetA2 | 10.10.131.0/24 |

以下の各文について、正しい場合は 「はい」を選択してください。そうでない場合は 「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

| Statements                                                            | Yes                   | No                    |
|-----------------------------------------------------------------------|-----------------------|-----------------------|
| A Site-to-Site connection can be established between VNET1 and VNET2. | <input type="radio"/> | <input type="radio"/> |
| VNET1 and VNET2 can be peered.                                        | <input type="radio"/> | <input type="radio"/> |
| VNET1 and VNETA can be peered.                                        | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

| Statements                                                            | Yes                              | No                               |
|-----------------------------------------------------------------------|----------------------------------|----------------------------------|
| A Site-to-Site connection can be established between VNET1 and VNET2. | <input checked="" type="radio"/> | <input type="radio"/>            |
| VNET1 and VNET2 can be peered.                                        | <input checked="" type="radio"/> | <input type="radio"/>            |
| VNET1 and VNETA can be peered.                                        | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

Answer Area

Statements

Yes

No

A Site-to-Site connection can be established between VNET1 and VNET2.

☒

☐

VNET1 and VNET2 can be peered.

☒

☐

VNET1 and VNETA can be peered.

☐

☒

In Microsoft Azure, connectivity between virtual networks depends on IP address space, region, and subscription constraints.

1## Site-to-Site VPN between VNET1 and VNET2 - YES:

Azure supports a Site-to-Site VPN connection between two VNets in different regions or subscriptions as long as both have non-overlapping address spaces and contain a VPN gateway. In this case, VNET1 (10.10.10.0/24 - West Europe) and VNET2 (172.16.0.0/16 - West US) have unique address spaces, so a S2S connection is possible.

2## VNET1 and VNET2 Peering - YES:

VNet peering enables direct communication over the Microsoft backbone network with low latency and no gateway requirement. Global VNet Peering supports VNets in different Azure regions (e.g., West Europe and West US) as long as address spaces do not overlap-which is true here.

3## VNET1 and VNETA Peering - NO:

VNETA's address space (10.10.128.0/17) overlaps with VNET1 (10.10.10.0/24), because both belong to the same 10.10.0.0/8 range. Azure explicitly blocks VNet peering between VNets with overlapping address spaces to avoid routing conflicts.

According to Microsoft Learn ("Virtual network peering - requirements and constraints"), peering is only supported when address spaces do not overlap. Therefore, VNET1 and VNETA cannot be peered.

質問: 107

オンプレミス ネットワークがあります。

VNet1 と VNet2 という 2 つの仮想ネットワークを含む Azure サブスクリプションがあります。VNet1 は IP アドレス空間 192.168.8.0/24 を使用します。VNet2 は IP アドレス空間 192.168.9.0/24 を使用します。

仮想ネットワークを構成する必要があります。ソリューションは以下の要件を満たす必要があります。


- \* VNet1 上のリソースが VNet2 上のリソースと通信できることを確認します。
- \* オンプレミス ネットワーク上のリソースが Azure リソースと通信できることを確認します。
- \* コストを最小限に抑えます。



正解:



Explanation:



To allow VNet1 (192.168.8.0/24) and VNet2 (192.168.9.0/24) to communicate, the lowest-overhead Azure- native option is virtual network peering. Azure documentation describes VNet peering as a way to connect two VNets so that resources can communicate using private IP addresses as if they were on the same network. Because the address spaces do not overlap, peering is supported and does not require deploying extra appliances. This also minimizes administrative effort compared to routing through gateways or NVAs. For on-premises to Azure connectivity while minimizing cost, the appropriate service is an Azure VPN Gateway with a site-to-site (S2S) VPN. Azure's guidance for hybrid connectivity positions VPN Gateway as the cost-effective encrypted tunnel over the public internet for connecting on-premises networks to Azure VNets. The alternative shown, ExpressRoute, provides private dedicated connectivity but is designed for higher throughput/enterprise requirements and generally increases cost due to circuit provisioning and associated charges. Therefore, combining VNet peering (VNet-to-VNet) with VPN Gateway (on-prem-to- Azure) satisfies both communication requirements and the cost constraint.

質問: 108

次の表のリソースを含む Azure サブスクリプションがあります。

| Name | Type                      |
|------|---------------------------|
| VM1  | Virtual machine           |
| VM2  | Virtual machine           |
| LB1  | Load balancer (Basic SKU) |

VM1 および VM2 に Web サーバー (IIS) サーバー ロールをインストールし、VM1 および VM2 を LB1 に追加します。  
LB1 は、LB1 の展示に示されているように構成されています。(LB1 タブをクリックします。)

Essentials ^

Resource group (change)  
VMRG

Location  
West Europe

Subscription name (change)  
Azure Pass

Subscription ID  
e66d2b22-fde8-4af2-9323-d43516f6eb4e

SKU  
Basic

Backend pool  
Backend1 (2 virtual machines)

Health probe  
Probe1 (HTTP:80/Probe1.htm)

Load balancing rule  
Rule1 (TCP/80)

NAT rules  
-

Public IP address  
104.40.178.194 (LB1)

Rule1 は、Rule1 の図に示すように構成されています ([Rule1] タブをクリックします)。

\* Name  
Rule1

\* IP Version  
☒ IPv4 ☐ IPv6

\* Frontend IP address ⓘ  
104.40.178.194 (LoadBalancerFrontEnd) ▼

Protocol  
☒ TCP ☐ UDP

\* Port  
80

\* Backend port ⓘ  
80

Backend pool ⓘ  
Backend1 (2 virtual machines) ▼

Health probe ⓘ  
Probe1 (HTTP:80/Probe1.htm) ▼

Session persistence ⓘ  
None ▼

Idle timeout (minutes) ⓘ

Floating IP (direct server return) ⓘ  
Disabled

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

**Answer Area**

| Statements                                                                                             | Yes                   | No                    |
|--------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| VM1 is in the same availability set as VM2.                                                            | <input type="radio"/> | <input type="radio"/> |
| If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2. | <input type="radio"/> | <input type="radio"/> |
| If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.          | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

| Statements                                                                                             | Yes                   | No                    |
|--------------------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| VM1 is in the same availability set as VM2.                                                            | <input type="radio"/> | <input type="radio"/> |
| If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2. | <input type="radio"/> | <input type="radio"/> |
| If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.          | <input type="radio"/> | <input type="radio"/> |

Explanation:

Answer Area

| Statements                                                                                             | Yes                              | No                               |
|--------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| VM1 is in the same availability set as VM2.                                                            | <input checked="" type="radio"/> | <input type="radio"/>            |
| If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2. | <input checked="" type="radio"/> | <input type="radio"/>            |
| If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.          | <input type="radio"/>            | <input checked="" type="radio"/> |

1## VM1 is in the same availability set as VM2 # Yes

In Azure, for a Basic SKU Load Balancer, the backend virtual machines must be in the same availability set (or in the same virtual machine scale set).

Basic Load Balancers cannot span across multiple availability sets or availability zones.

The screenshot shows that LB1 uses the Basic SKU, so by requirement, VM1 and VM2 must be in the same availability set.

Therefore, this statement is true.

2## If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2 # Yes A load balancing rule defines how traffic is distributed to the backend pool.

The rule "Rule1" uses TCP port 80 for both frontend and backend.

A health probe (Probe1.htm) monitors the health of backend VMs.

When both VMs respond successfully to Probe1.htm, the load balancer distributes TCP traffic on port 80 evenly between the two healthy VMs.

Therefore, when Probe1.htm is available and healthy on both VMs, LB1 distributes HTTP traffic (port 80) between VM1 and VM2.

This statement is true.

3## If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all ports # No Load balancers in Azure do not automatically balance traffic unless a load-balancing rule exists.

Each load balancing rule explicitly defines which ports and protocols are balanced.

If you delete Rule1, there will be no configuration telling the load balancer which ports to listen to or where to send traffic.

The load balancer will not distribute any traffic automatically for other ports.

Therefore, deleting the rule disables traffic distribution entirely, and this statement is false.

Microsoft Azure Administrator Official Documentation Extract (AZ-104 Reference):

"Basic Load Balancer can only balance traffic among VMs within a single availability set or virtual machine scale set."

"A load-balancing rule defines how traffic is distributed to the backend pool. Deleting a rule stops load balancing for that port."

"Health probes determine which backend pool instances receive new flows." (Source: Microsoft Learn - Azure Load Balancer Overview, Create and Manage Load Balancer Rules, Health Probes in Azure Load Balancer)

質問: 109

Azure Resource Manager テンプレートを使用して、仮想マシン スケール セットに Windows Server 2022 を実行する複数の Azure 仮想マシンをデプロイする予定です。

仮想マシンをデプロイした後、すべての仮想マシンで NGINX が使用可能であることを確認する必要があります。

何を使うべきでしょうか？

A. Microsoft Intune デバイス構成プロファイル

B. Microsoft エントラ アプリケーション プロキシ

C. Azure カスタム スクリプト拡張機能

D. Azure App Service の Department Center

正解: (正解を表示します)

When deploying Azure Virtual Machines (VMs) - individually or within a Virtual Machine Scale Set (VMSS) - administrators often need to automate post-deployment configuration tasks such as installing software, running configuration scripts, or setting up environment dependencies.

According to Microsoft Azure Administrator documentation, the Azure Custom Script Extension is designed precisely for this purpose. It downloads and executes scripts on Azure VMs after deployment. The extension integrates directly with Azure Resource Manager (ARM) templates, Azure CLI, or PowerShell, enabling automated installation and configuration of software packages like NGINX, IIS, or custom utilities.


The Custom Script Extension runs PowerShell scripts on Windows VMs and shell scripts on Linux VMs. In the case of deploying Windows Server 2022 VMs in a scale set, you can embed the Custom Script Extension JSON configuration into your ARM template, specifying the script's location in Azure Storage or GitHub and the command to execute (for example, Install-WindowsFeature Web-Server or an NGINX installer command).

This approach aligns with Microsoft's best practices for Azure automation and infrastructure-as-code, ensuring consistent, repeatable deployments across all instances in the scale set. It also minimizes manual intervention, reduces configuration drift, and ensures each VM instance has the required software components available immediately after provisioning.

質問: 110

次の図に示すように構成される VM1 という名前の Azure 仮想マシンを作成する予定です。

VM1 の計画されたディスク構成を次の図に示します。



Instance details

Virtual machine name \*

Region \*

Availability options

Image \*   
Browse all public and private images

Azure Spot instance ☐ Yes ☒ No

Size \*   
1 vcpu, 3.5 GiB memory (ZAR 632.47/month)  
[Change size](#)

The planned disk configurations for VM1 are shown in the following exhibit.

Basics Disks Networking Management Advanced Tags Review + create

Azure VMs have one operating system disk and a temporary disk for short-term storage. You can attach additional data disks. The size of the VM determines the type of storage you can use and the number of data disks allowed. [Learn more](#)

Disk options

OS disk type \*

The selected VM size supports premium disks. We recommend Premium SSD for high IOPS workloads. Virtual machines with Premium SSD disks qualify for the 99.9% connectivity SLA.

Enable Ultra Disk compatibility ☐ Yes ☒ No  
Ultra Disks are only available when using Managed Disks.

Data disks

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

**i** Adding unmanaged data disks is currently not supported at the time of VM creation. You can add them after the VM is created.

Advanced

Use managed disks ☒ No ☐ Yes

Storage account \*   
Create new

VM1 をアベイラビリティーゾーンに作成できることを確認する必要があります。

どの 2 つの設定を変更する必要がありますか？それぞれの正解は解決策の一部を示しています。

注意: 正しい選択ごとに 1 ポイントが加算されます。

- A. マネージドディスクを使用する
- B. 可用性オプション
- C. OSディスクの種類

D. サイズ

E. 画像

正解: A,B (コメントを發表する)

https://docs.microsoft.com/en-us/azure/site-recovery/move-azure-vms-avset-azone https://docs.microsoft.com/en-us/azure/virtual-machines/windows/create-portal-availability-zone https://docs.microsoft.com/en-us/azure/virtual-machines/manage-availability https://docs.microsoft.com/en-us/azure/availability-zones/az-overview#availability-zones

質問: 111

ストレージアカウントのコンテンツに計画されている変更を適用する必要があります。コンテンツを整理するために、どのコンテナとファイル共有を使用できますか？

A. share1のみ

B. cont1 と share1 のみ

C. share1 と share2 のみ

D. cont1、share1、share2のみ

E. cont1、cont2、share1、share2

正解: (正解を表示します)

In the scenario, storage1 is configured as StorageV2 with Hierarchical namespace = Yes, while storage2 is configured as StorageV2 with Hierarchical namespace = No.

From Microsoft's Azure Storage Documentation and AZ-104 Study Guide, the following principles apply:

- \* A hierarchical namespace (enabled when the storage account has Azure Data Lake Storage Gen2 capabilities) allows the use of directories within containers to organize data.
- \* The hierarchical namespace provides directory and file-level structure similar to a file system. This is supported only for blob containers, not for Azure Files.
- \* Azure Files (file shares) do not depend on hierarchical namespaces and cannot have directories in the same way Data Lake Gen2 does - directories can exist inside the share but not in the blob container sense.
- \* The planned change states that you must use directories whenever possible to organize content.

Therefore, only storage accounts with hierarchical namespace enabled can use directory structures - that's storage1.

In this case:

- \* storage1 (Hierarchical namespace = Yes) # supports containers (like cont1) and file shares (like share1).
- \* storage2 (Hierarchical namespace = No) # does not support directories within blob containers (Data Lake structure).

Hence, you can use only cont1 (container in storage1) and share1 (file share in storage1) to organize content as required.

This is directly supported by the Microsoft documentation on Data Lake Storage Gen2:

"When you enable the hierarchical namespace for a storage account, you can organize objects into directories and subdirectories. This capability is available only for accounts configured for Data Lake Storage Gen2." Final

Verified Answer: # B. cont1 and share1 only

質問: 112

次の表のリソースを含む Azure サブスクリプションがあります。

| Name  | Type                               |
|-------|------------------------------------|
| VMRG  | Resource group                     |
| VNet1 | Virtual network                    |
| VNet2 | Virtual network                    |
| VM5   | Virtual machine connected to VNet1 |
| VM6   | Virtual machine connected to VNet2 |

Azure で、adatum.com という名前のプライベート DNS ゾーンを作成し、VNet2 に仮想ネットワーク リンクを追加して、自動登録を有効にします。

adatum.com ゾーンは、次の図に示すように構成されています。

Resource group (change)  
vmrg

Subscription (change)  
Azure Pass

Subscription ID  
a4fde29b-d56a-4f6c-8298-6c53cd0b720c

Name server 1  
-

Name server 2  
-

Name server 3  
-

Name server 4  
-

Tags (change)  
Click here to add tags

| NAME | TYPE | TTL  | VALUE                                                                                                                                                             |
|------|------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| @    | SOA  | 3600 | Email: azuredns-hostmaster.microsoft.com<br>Host: internal.cloudapp.net<br>Refresh: 3600<br>Retry: 300<br>Expire: 2419200<br>Minimum TTL: 300<br>Serial number: 1 |
| vm1  | A    | 3600 | 10.1.0.4                                                                                                                                                          |
| vm9  | A    | 3600 | 10.1.0.12                                                                                                                                                         |

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注: 正解ごとに1ポイント獲得

Answer Area

| Statements                                                                    | Yes                   | No                    |
|-------------------------------------------------------------------------------|-----------------------|-----------------------|
| The A record for VM5 will be registered automatically in the adatum.com zone. | <input type="radio"/> | <input type="radio"/> |
| VM5 can resolve VM9.adatum.com.                                               | <input type="radio"/> | <input type="radio"/> |
| VM6 can resolve VM9.adatum.com.                                               | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

| Statements                                                                    | Yes                              | No                               |
|-------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| The A record for VM5 will be registered automatically in the adatum.com zone. | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM5 can resolve VM9.adatum.com.                                               | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM6 can resolve VM9.adatum.com.                                               | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

Answer Area



| Statements                                                                    | Yes                              | No                               |
|-------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| The A record for VM5 will be registered automatically in the adatum.com zone. | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM5 can resolve VM9.adatum.com.                                               | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM6 can resolve VM9.adatum.com.                                               | <input checked="" type="radio"/> | <input type="radio"/>            |

質問: 113

Microsoft Entra テナントがあり、External User という名前のユーザーが含まれています。External User は、external95@gmail.com を使用してテナントに対して認証します。外部ユーザーが contractor@gmail.com を使用してテナントに対して認証していることを確認する必要があります。

概要ブレードから構成する必要がある 2 つの設定はどれですか? 回答するには、回答領域で適切な設定を選択します。注: 正解ごとに 1 ポイントが加算されます。

Answer Area

External User

User

Edit properties

Delete

Refresh

Reset password

Revoke sessions

Manage view

Got feedback?

Overview

Audit logs

Sign-in logs

Diagnose and solve problems

Manage

Custom security attributes (preview)

Assigned roles

Administrative units

Groups

Applications

Licenses

Devices

Azure role assignments

Authentication methods

Troubleshooting + Support

New support request

Overview

Monitoring

Properties

Basic info

EU

External User

external195\_gmail.com#EXT#@sk230415outlook.onmicrosoft.com

Guest

User principal name

external195\_gmail.com#EXT#@sk230415outlook.onmicroso...

Group membe...

0

Object ID

2b353249-fa3d-4c0e-b69d-fa6c6c60fa1c

Applications

0

Created date time

Apr 30, 2023, 11:58 AM

Assigned roles

0

User type

Guest

Assigned lice

0

Identities

mail

30, 2023, 11:58 AM

My Feed

Account status

Enabled

Edit

Sign-ins

Last sign-in: -- --

See all sign-ins

B2B collaboration

Invitation state: Accepted

Reset redemption status

正解:

External User

User

Search

Overview

Audit logs

Sign-in logs

Diagnose and solve problems

Manage

Custom security attributes (preview)

Assigned roles

Administrative units

Groups

Applications

Licenses

Devices

Azure role assignments

Authentication methods

Troubleshooting + Support

New support request

Edit properties

Delete

Refresh

Reset password

Revoke sessions

Manage view

Got feedback?

Overview

Monitoring

Properties

Basic info

EU

External User

external195\_gmail.com#EXT#@sk230415outlook.onmicrosoft.com

Guest

User principal name

external195\_gmail.com#EXT#@sk230415outlook.onmicroso...

Object ID

2b353249-fa3d-4c8e-b69d-fa6c6c60fa1c

Created date time

Apr 30, 2023, 11:58 AM

User type

Guest

Identities

mail 30, 2023, 11:58 AM

Group members

0

Applications

0

Assigned roles

0

Assigned licenses

0

My Feed

Account status

Enabled

Edit

Sign-ins

Last sign-in: -- --


See all sign-ins

B2B collaboration

Invitation state: Accepted

Reset redemption status

Explanation:



In Microsoft Entra ID (formerly Azure Active Directory), guest users (External Users) are typically added via B2B collaboration invitations. They authenticate using the identity provider (IdP) associated with the email they were invited with - such as a Microsoft account, Gmail (Google), or another organization's Azure AD tenant.

When you need to change the authentication method or identity for a guest user - for example, from external195@gmail.com to contractor@gmail.com - you must modify their identities and ensure that their account status is enabled and synchronized with the new sign-in identity.

Step-by-step from Microsoft Documentation:

From Microsoft Learn - "Manage external collaboration settings in Microsoft Entra ID" and "Manage guest accounts in Azure AD":

#### Identities

Under the Overview blade of the user, the Identities section lists all linked sign-in identities for that account.

You can remove the old identity (e.g., external195@gmail.com) and add the new one (contractor@gmail.com) as an alternate identity or replace the existing one.

This determines which external account the guest uses to authenticate to your tenant.

#### Account status

Once the new identity is added, ensure the Account status remains Enabled so that the user can sign in successfully.

If disabled, the user cannot authenticate even with the new email address.

This setting controls whether the guest user object is active in the directory.

Other fields such as User principal name, Object ID, or Reset password cannot directly modify the external identity since guest authentication relies on the federated identity provider of the external email domain.

Therefore, the administrator must use Identities to register or replace the authentication address, and Account status to ensure the user can log in using the new account.

質問: 114

Azure Blob Storage と Azure File Storage を使用する storage1 という名前の Azure Storage アカウントがあります。  
AzCopy を使用して、storage1 の BLOB ストレージとファイル ストレージにデータをコピーする必要があります。  
各タイプのストレージにはどの認証方法を使用する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

Blob storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

File storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

正解:

Blob storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

File storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

Explanation:

Blob storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

File storage:

Azure Active Directory (Azure AD) only

Shared access signatures (SAS) only

Access keys and shared access signatures (SAS) only

Azure Active Directory (Azure AD) and shared access signatures (SAS) only

Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

You can provide authorization credentials by using Azure Active Directory (AD), or by using a Shared Access Signature (SAS) token.

Box 1:

Both Azure Active Directory (AD) and Shared Access Signature (SAS) token are supported for Blob storage.

Box 2:

Only Shared Access Signature (SAS) token is supported for File storage.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10>

質問: 115

contoso.com という名前の登録済み DNS ドメインがあります。

contoso.com という名前のパブリック Azure DNS ゾーンを作成します。

contoso.com ゾーンで作成されたレコードがインターネットから解決可能であることを確認する必要があります。

何をすべきでしょうか？

- A. DNS ドメイン レジストラの NS レコードを変更します。
- B. contoso.com に NS レコードを作成します。
- C. contoso.com に SOA レコードを作成します。
- D. DNS ドメイン レジストラの SOA レコードを変更します。

正解: (正解を表示します)

When a public Azure DNS zone is created, Azure automatically assigns a set of authoritative name servers (NS records) to the zone. However, simply creating the DNS zone in Azure does not make it resolvable from the internet. For external resolution to work, the domain registrar must delegate authority to Azure DNS.

Microsoft Azure documentation clearly states that to make DNS records resolvable publicly, you must update the domain registrar's NS records to point to the Azure-assigned name servers. This action establishes Azure DNS as the authoritative DNS provider for the domain.

Creating NS or SOA records within the Azure DNS zone itself does not affect external resolution unless the registrar delegation is completed. The SOA record is automatically created and managed by Azure DNS and must not be modified at the registrar.

質問: 116

VNet1 という名前の Azure 仮想ネットワークがあり、サイト間 VPN を使用してオンプレミス ネットワークに接続しています。VMet1 には、Subnet1 という名前のサブネットが 1 つ含まれています。

Subnet1 は、NSG1 という名前のネットワーク セキュリティ グループ (NSG) に関連付けられています。Subnet1 には、ILB1 という名前の基本的な内部ロード バランサーが含まれています。ILB1 のバックエンド プールには 3 台の Azure 仮想マシンがあります。

ILB1 に接続する IP アドレスに関するデータを収集する必要があります。収集したデータに対して、Azure ポータルから対話型クエリを実行できる必要があります。

どうすればいいのでしょうか？ 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Resource to create:

An Azure Event Grid

An Azure Log Analytics workspace

An Azure Storage account

Resource on which to enable diagnostics:

ILB1

NSG1

The Azure virtual machines

正解:

Resource to create:

An Azure Event Grid

An Azure Log Analytics workspace

An Azure Storage account

Resource on which to enable diagnostics:

ILB1

NSG1

The Azure virtual machines

Explanation:

Box 1: An Azure Log Analytics workspace

In the Azure portal you can set up a Log Analytics workspace, which is a unique Log Analytics environment with its own data repository, data sources, and solutions.

Box 2: NSG1

NSG flow logs allow viewing information about ingress and egress IP traffic through a Network security group. Through this, the IP addresses that connect to the ILB can be monitored when the diagnostics are enabled on a Network Security Group.

We cannot enable diagnostics on an internal load balancer to check for the IP addresses.

As for Internal LB, it is basic one. Basic can only connect to storage account. Also, Basic LB has only activity logs, which doesn ' t include the connectivity workflow. So, we need to use NSG to meet the mentioned requirements.

質問: 117

Azure サブスクリプションをお持ちです。次の表に示すリソースをデプロイする予定です。

| Name  | Type                                    |
|-------|-----------------------------------------|
| IP1   | Microsoft.Network/publicIPAddresses     |
| NSG1  | Microsoft.Network/networkSecurityGroups |
| VNET1 | Microsoft.Network/virtualNetworks       |
| NIC1  | Microsoft.Network/networkInterfaces     |
| VM1   | Microsoft.Compute/virtualMachines       |

リソースのデプロイに使用する単一のAzure Resource Manager (ARM) テンプレートを作成する必要があります。VM1のdependsOnセクションに追加する必要があるリソースはどれですか？

- A. IP1
- B. VNET1
- C. NIC1
- D. NSG1

正解: [\(正解を表示します\)](#)

When deploying a virtual machine via an ARM template, dependencies must ensure resources are created in the proper order. A VM requires a network interface (NIC) at creation time; therefore, the VM resource must depend on the NIC.

The NIC itself depends on the virtual network, subnet, network security group, and public IP (if applicable).

Microsoft ARM template guidance specifies that the VM's dependsOn should reference the NIC, allowing Azure Resource Manager to implicitly handle the remaining network dependencies.

質問: 118

Sub1 と Sub2 という名前の 2 つの Azure サブスクリプションがあり、それぞれが別々の Microsoft Entra テナントにリンクされています。

次の表に示す仮想ネットワークがあります。

| Name  | Location   | Subscription |
|-------|------------|--------------|
| VNet1 | East US    | Sub1         |
| VNet2 | East US    | Sub1         |
| VNet3 | West US    | Sub1         |
| VNet4 | East US    | Sub2         |
| VNet5 | Central US | Sub2         |

Sub1 と Sub2 という名前の 2 つの Azure サブスクリプションがあり、それぞれが別々の Microsoft Entra テナントにリンクされています。

次の表に示す仮想ネットワークがあります。

VNet1 とピアリングできる仮想ネットワークはどれですか？

- A. VNet2 と VNet4 のみ
- B. VNet2のみ
- C. VNet2、VNet3、VNet4、および VNet5
- D. VNet2、VNet3、および VNet4 のみ
- E. VNet2 と VNet3 のみ

正解: [\(正解を表示します\)](#)

質問: 119

次の表に示すコンテナを含む Azure サブスクリプションがあります。

| Name      | Type                    |
|-----------|-------------------------|
| Backup1   | Backup vault            |
| Recovery1 | Recovery Services vault |

次の表に示すリソースを含むストレージ アカウントを作成します。

| Name   | Type           |
|--------|----------------|
| cont1  | Blob container |
| share1 | File share     |

cont1とshare1をどの保管庫にバックアップできますか？回答するには、回答エリアで適切な選択肢を選択してください。注：正解は1点です。

Answer Area

cont1:

Backup1 only

Backup1 only

Recovery1 only

Backup1 or Recovery1

Cannot be backed up to Backup1 or Recovery1

share1:

Recovery1 only

Backup1 only

Recovery1 only

Backup1 or Recovery1

Cannot be backed up to Backup1 or Recovery1

正解:

cont1:

Backup1 only

Backup1 only

Recovery1 only

Backup1 or Recovery1

Cannot be backed up to Backup1 or Recovery1

share1:

Recovery1 only

Backup1 only

Recovery1 only

Backup1 or Recovery1

Cannot be backed up to Backup1 or Recovery1

Explanation:

Answer Area



In Microsoft Azure, backup vaults and Recovery Services vaults serve distinct purposes depending on the type of data source being protected. The Azure Administrator study guide and Microsoft documentation outline the following key differences and supported workloads:

- \* Backup Vault (Azure Backup Center - Modern Vault):

- \* The Backup vault is designed primarily for Azure Blob storage backup.

- \* It supports backing up Azure Blobs (block blobs) in a storage account.

- \* You can configure backup policies directly in Backup Center and use the "Azure Blobs backup" feature.

- \* This vault type is managed using the Azure Resource Manager (ARM) model and supports new backup scenarios introduced after 2021.

- \* Hence, in this question, cont1 (a blob container) can be backed up only to Backup1 (Backup vault).

- \* Recovery Services Vault:

- \* The Recovery Services vault supports backups for Azure virtual machines (VMs), Azure Files shares, on-premises workloads using the Microsoft Azure Recovery Services (MARS) agent, and SQL/SAP HANA databases on Azure VMs.

- \* It does not support Blob container backup but does support Azure file share backup.

- \* Therefore, share1 (a file share) must be backed up using Recovery1 (Recovery Services vault).

The Azure Backup architecture documentation emphasizes that these vaults are distinct and not interchangeable:

"Recovery Services vaults are used to back up Azure VMs, Azure Files, and on-premises workloads. Backup vaults are used to protect Azure Blobs. You cannot back up Blobs to a Recovery Services vault or back up Azure Files to a Backup vault." Thus, following Microsoft's official storage backup configuration:

- \* cont1 (Blob container) # Backup1 only

- \* share1 (File share) # Recovery1 only

Final Verified Answer:

# cont1 # Backup1 only

# share1 # Recovery1 only

質問: 120

contoso.onmicrosoft.com という名前の Azure Active Directory (Azure AD) テナントがあります。

ユーザー管理者ロールは、Admin1 という名前のユーザーに割り当てられます。

外部パートナーには、user1@outlook.com サインインを使用する Microsoft アカウントがあります。

Admin1 は、外部パートナーを Azure AD テナントにサインインするように招待しようとしたのですが、次のエラー メッセージが表示されました: ユーザー user1@outlook.com を招待できません - 一般的な承認例外です。」

Admin1 が外部パートナーを Azure AD テナントにサインインするように招待できることを確認する必要があります。

何をすべきでしょうか？

A. [ロールと管理者] ブレードから、セキュリティ管理者ロールを Admin1 に割り当てます。

B. [組織関係] ブレードから、ID プロバイダーを追加します。

C. [カスタム ドメイン名] ブレードから、カスタム ドメインを追加します。

D. ユーザー設定ブレードから、外部コラボレーション設定を変更します。

正解: D ([コメントを发表する](#))

In Azure AD, the ability to invite external (guest) users is controlled through External collaboration settings, found under Users # External collaboration settings in the Azure AD portal.

If a User administrator attempts to invite an external user but receives a "Generic authorization exception", this typically means the organization's external collaboration restrictions prevent user invitations.

According to Microsoft documentation (Configure external collaboration settings in Azure AD):

"To enable administrators or users to invite external guests, you must configure the External collaboration settings. You can define who can invite guest users into the directory, including administrators, users, or only specific roles." By updating these settings to allow User administrators to send invitations, Admin1 will be able to successfully invite the external partner (user1@outlook.com).

質問: 121

次の表に示す仮想ネットワークを含む Azure サブスクリプションがあります。

| Name  | Subnet  | Subnet-associated network security group (NSG) | Peered with |
|-------|---------|------------------------------------------------|-------------|
| VNet1 | Subnet1 | NSG1                                           | VNet2       |
| VNet2 | Subnet2 | NSG2                                           | VNet1       |

サブスクリプションには、次の表に示す仮想マシンが含まれています。

| Name | Connected to |
|------|--------------|
| VM1  | Subnet1      |
| VM2  | Subnet2      |

サブスクリプションには、次の表に示す Azure App Service Web アプリが含まれています。

| Name    | Description                                                                  |
|---------|------------------------------------------------------------------------------|
| WebApp1 | Uses the Premium pricing tier and has virtual network integration with VNet1 |
| WebApp2 | Uses the Isolated pricing tier and is deployed to Subnet2                    |

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Microsoft

| Statements                                | Yes                   | No                    |
|-------------------------------------------|-----------------------|-----------------------|
| WebApp1 can communicate with VM2.         | <input type="radio"/> | <input type="radio"/> |
| NSG1 controls inbound traffic to WebApp1. | <input type="radio"/> | <input type="radio"/> |
| WebApp2 can communicate with VM1.         | <input type="radio"/> | <input type="radio"/> |

正解:

Microsoft

Answer Area

| Statements                                | Yes                              | No                               |
|-------------------------------------------|----------------------------------|----------------------------------|
| WebApp1 can communicate with VM2.         | <input type="radio"/>            | <input checked="" type="radio"/> |
| NSG1 controls inbound traffic to WebApp1. | <input type="radio"/>            | <input checked="" type="radio"/> |
| WebApp2 can communicate with VM1.         | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

WebApp1 can communicate with VM2. No, this is not correct. According to the tables, WebApp1 is integrated with VNet1, which has a peering connection with VNet2. However, VM2 is in VNet3, which is not peered with VNet1 or VNet2. Therefore, WebApp1 cannot communicate with VM2 across different virtual networks1.

NSG1 controls inbound traffic to WebApp1. No, this is not correct. According to the tables, NSG1 is associated with Subnet1 in VNet1, which is integrated with WebApp1. However, network security groups only control outbound traffic from App Service apps to virtual networks, not inbound traffic to App Service apps from virtual networks2. Therefore, NSG1 does not control inbound traffic to WebApp1. WebApp2 can communicate with VM1. Yes, this is correct. According to the tables, WebApp2 is integrated with VNet3, which has a peering connection with VNet2. VM1 is in Subnet2 in VNet2, which has a network security group named NSG2 that allows inbound traffic from any source on port 803. Therefore, WebApp2 can communicate with VM1 on port 80 across peered virtual networks.

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%**ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 122

storage1 という Azure Storage アカウントがあり、その中に containers という BLOB コンテナが含まれています。container! に追加された新しいコンテンツが 1 年間変更されないようにする必要があります。どのような設定が必要ですか？

- A. アクセスレベル
- B. アクセス制御（午前時）設定
- C. アクセスポリシー
- D. アクセス層

正解: [\(正解を表示します\)](#)

To prevent new content added to an Azure Blob container from being modified or deleted for a specific duration (in this case, one year), you must configure an immutability policy by setting a time-based retention policy or legal hold within an access policy on the container.

According to Microsoft's Azure Storage documentation, immutability policies are configured under

"Immutable blob storage", which allows you to store data in a WORM (Write Once, Read Many) state. When a time-based retention policy is set (for example, one year), any blob data added to that container cannot be modified or deleted until the retention period expires.

The configuration is done by defining an access policy on the container and specifying parameters like retentionPeriodInDays. This ensures compliance with regulations such as SEC 17a-4(f), CFTC 1.31(d), and FINRA Rule 4511.

Other options such as access level, IAM settings, and access tier control visibility, permissions, or storage costs but do not enforce immutability or write protection.

Hence, to achieve the goal of preventing modifications to new blobs for one year, you must configure an access policy with a time-based immutability (retention) policy on the container.

# Final Verified Answer: C. an access policy

質問: 123

Azure サブスクリプションをお持ちです。

Azure Resource Manager テンプレートを使用して、Azure Bastion を使用する VNET1 という名前の仮想ネットワークをデプロイする予定です。

テンプレートをどのように完成させるべきですか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

## Answer Area

```
{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1",
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"]
    },
    "subnets": [
      {
        "name": "
          AzureBastionSubnet
          AzureFirewallSubnet
          LAN01
          RemoteAccessSubnet
        ",
        "properties": {
          "addressPrefix": "
            10.10.10.0/27
            10.10.10.0/29
            10.10.10.0/30
          "
        }
      },
      {
        "name": "LAN02",
        "properties": {
          "addressPrefix": "10.10.10.128/25"
        }
      }
    ]
  }
}
```

正解:

```
{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1",
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"]
    },
    "subnets": [
      {
        "name": 

- AzureBastionSubnet
- AzureFirewallSubnet
- LAN01
- RemoteAccessSubnet


        "properties": {
          "addressPrefix": 

- 10.10.10.0/27
- 10.10.10.0/29
- 10.10.10.0/30


        },
      },
      {
        "name": "LAN02",
        "properties": {
          "addressPrefix": "10.10.10.128/25"
        }
      }
    ]
  }
}
```

Explanation:

```

{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1"
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"]
    },
    "subnets": [
      {
        "name":
        "properties": {
          "addressPrefix":

```

▼

- AzureBastionSubnet
- AzureFirewallSubnet
- LAN01
- RemoteAccessSubnet

▼

- 10.10.10.0/27
- 10.10.10.0/29
- 10.10.10.0/30

```

      }
    ],
  },
}

```

When deploying Azure Bastion using an Azure Resource Manager (ARM) template, Microsoft requires that the subnet dedicated to Bastion follows strict configuration guidelines. According to the Microsoft Azure Administrator documentation, the Azure Bastion service must be deployed within a specifically named subnet called "AzureBastionSubnet".

This subnet name is case-sensitive and mandatory - any deviation from this name (e.g., "BastionSubnet" or "AzureBastion") will cause deployment failure.

Additionally, the subnet must have a minimum prefix of /27 to provide sufficient IP addresses for the Bastion- managed instances and scaling operations. Subnets smaller than /27 (such as /28 or /29) are not supported because Bastion requires multiple internal IP addresses for operations like scaling, management, and secure connectivity.

In the ARM template structure:

```

{
  " type " : " Microsoft.Network/virtualNetworks " ,
  " name " : " VNET1 " ,
  " apiVersion " : " 2019-02-01 " ,
  " location " : " [resourceGroup().location] " ,
  " properties " : {
    " addressSpace " : {

```

```
" addressPrefixes " : [ " 10.10.10.0/24 " ]
},
" subnets " : [
{
" name " : " AzureBastionSubnet " ,
" properties " : {
" addressPrefix " : " 10.10.10.0/27 "
}
},
{
" name " : " LAN02 " ,
" properties " : {
" addressPrefix " : " 10.10.10.128/25 "
}
}
]
}
```

This ensures that Azure Bastion can be deployed properly within the virtual network VNET1, providing secure RDP/SSH connectivity to VMs without exposing public IPs.

Extract from Microsoft Documentation:

"Azure Bastion must be deployed in a dedicated subnet named 'AzureBastionSubnet'. The subnet must be at least /27 or larger (/26, /25, etc.). This subnet is used exclusively by the Bastion host and cannot contain other resources." (Source: Microsoft Learn - Deploy Azure Bastion using Azure Resource Manager templates, and Azure Bastion service configuration requirements)

質問: 124

次の表に示す Azure 仮想マシンがあります。

| Name | Azure region |
|------|--------------|
| VM1  | West Europe  |
| VM2  | West Europe  |
| VM3  | North Europe |
| VM4  | North Europe |

- VM1 と VM2 を保護する Recovery Services コンテナがあります。
- Recovery Services を使用して VM3 と VM4 を保護する必要があります。
- まず何をすべきでしょうか？
- A. 新しい Recovery Services コンテナを作成します。
  - B. VM3 および VM4 の拡張機能を構成します。
  - C. ストレージ アカウントを作成します。
  - D. 新しいバックアップ ポリシーを作成します。

正解: B (コメントを发表する)

Azure Backup uses the Recovery Services vault to protect virtual machines by creating restore points stored in Azure. Before a virtual machine can be backed up, it must have the Azure Backup extension (also called VMBackup extension) properly installed and configured on the VM.

According to Microsoft Learn - "Back up and restore encrypted virtual machines" and "Back up Azure VMs with the Azure Backup service", the backup process for a VM requires the following:

"When you enable backup for an Azure virtual machine, the Azure Backup service installs an extension on the VM to coordinate the snapshot and backup process. If the extension is missing or not installed properly, backups cannot be triggered for that VM." Explanation of the Scenario You already have a Recovery Services vault protecting VM1 and VM2.

You now need to add VM3 and VM4 to the same vault for protection.

Before they can be added, you must ensure that:

The VMs are running in the same region as the Recovery Services vault (because a vault can only protect resources within its region).

The VMBackup extension is installed and registered on each VM.

If the backup extension is missing or not configured correctly, Azure Backup cannot communicate with the VM to take snapshots or register it in the vault.

Once the extension is properly configured, you can then enable backup from the Recovery Services vault without needing to create a new vault, policy, or storage account - because all those resources already exist and can be reused.

Other Options Explained

- A). Create a new Recovery Services vault: # Not necessary unless the new VMs are in a different region from the existing vault.
- C). Create a storage account: # Azure Backup automatically manages backup storage within the vault; no separate storage account is required.
- D). Create a new backup policy: # Not needed unless you want a different backup schedule or retention. You can use the existing policy.

# Final Verified Answer:

- B). Configure the extensions for VM3 and VM4.

Microsoft Learn Extract (Supporting Documentation):

"To enable Azure Backup for virtual machines, the Azure Backup extension must be installed. This extension coordinates the snapshot and backup process and is installed automatically when backup is enabled."

"If the extension is not installed, you must configure it before enabling protection." (Source: Microsoft Learn - Back up Azure VMs with the Azure Backup service, Azure Recovery Services vault overview)

質問: 125


次の表に示す Azure 仮想マシンを含む Azure サブスクリプションがあります。

| Name | Operating system             | Subnet  | Virtual network |
|------|------------------------------|---------|-----------------|
| VM1  | Windows Server 2019          | Subnet1 | VNET1           |
| VM2  | Windows Server 2019          | Subnet2 | VNET1           |
| VM3  | Red Hat Enterprise Linux 7.7 | Subnet3 | VNET1           |

次の表に示す設定を使用するように仮想マシンのネットワークインターフェースを構成します。

| Name | DNS server    |
|------|---------------|
| VM1  | None          |
| VM2  | 192.168.10.15 |
| VM3  | 192.168.10.15 |

VNET1 の設定から、次の図に示す DNS サーバーを構成します。



仮想マシンは、IP アドレス 192.168.10.15 を持つ DNS サーバーと、IP アドレス 193.77.134.10 を持つ DNS サーバーに正常に接続できます。

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

|                                                | Yes                   | No                    |
|------------------------------------------------|-----------------------|-----------------------|
| VM1 connects to 193.77.134.10 for DNS queries. | <input type="radio"/> | <input type="radio"/> |
| VM2 connects to 193.77.134.10 for DNS queries. | <input type="radio"/> | <input type="radio"/> |
| VM3 connects to 192.168.10.15 for DNS queries. | <input type="radio"/> | <input type="radio"/> |

正解:

|                                                | Yes                              | No                               |
|------------------------------------------------|----------------------------------|----------------------------------|
| VM1 connects to 193.77.134.10 for DNS queries. | <input checked="" type="radio"/> | <input type="radio"/>            |
| VM2 connects to 193.77.134.10 for DNS queries. | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM3 connects to 192.168.10.15 for DNS queries. | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

| Statements                                     | Yes                              | No                               |
|------------------------------------------------|----------------------------------|----------------------------------|
| VM1 connects to 193.77.134.10 for DNS queries. | <input checked="" type="radio"/> | <input type="radio"/>            |
| VM2 connects to 193.77.134.10 for DNS queries. | <input type="radio"/>            | <input checked="" type="radio"/> |
| VM3 connects to 192.168.10.15 for DNS queries. | <input checked="" type="radio"/> | <input type="radio"/>            |

In Azure, DNS resolution order follows a specific hierarchy based on custom DNS configurations at the virtual network (VNet) and network interface (NIC) levels. According to Microsoft Azure Administrator documentation, when you specify custom DNS servers for a virtual network, those DNS servers become the default for all virtual machines connected to that VNet - unless a specific VM's network interface has its own DNS settings configured. In that case, the NIC-level configuration overrides the VNet-level configuration. Here's how it applies to this scenario: VNET1 has been configured with a custom DNS server of 193.77.134.10 (visible in the VNet's DNS configuration screenshot). VM1 does not have a custom DNS server specified on its NIC, so it inherits the VNet-level DNS server (193.77.134.10). VM2 and VM3 both have explicit custom DNS settings pointing to 192.168.10.15 at the NIC level. Since NIC-level settings take precedence over the VNet's configuration, these VMs will use 192.168.10.15 for all DNS queries. This hierarchy is described in Azure documentation under: "DNS name resolution for virtual machines in Azure virtual networks," which explains that NIC-specific settings override VNet settings, and if no custom DNS is defined at either level, Azure's default internal DNS (168.63.129.16) is used. In summary: VM1 inherits DNS = 193.77.134.10 (Yes) VM2 uses custom DNS = 192.168.10.15 (No for 193.77.134.10) VM3 uses custom DNS = 192.168.10.15 (Yes) Thus, the correct and Microsoft-verified answers are:

# VM1: Yes, VM2: No, VM3: Yes

質問: 126

次の図に示すような Azure ポリシーがあります。

SCOPE

\* Scope (Learn more about setting the scope)

Subscription 1

Exclusions

Subscription 1/ContosoRG1

BASICS

\* Policy definition

Not allowed resource types

\* Assignment name ⓘ

Not allowed resource types

Assignment ID

/subscriptions/5eb8d0b6-ce3b-4ce0-a631-9f5321bedabb/providers/Microsoft.Authorization/policyAssignments/0e6fb866bf854f54accae2a9

Description

Assigned by

admin1@contoso.com

PARAMETERS

\* Not allowed resource types ⓘ

Microsoft.Sql/servers

この政策の効果は何ですか？

- A. Subscnption1 内の任意のリソース グループに Azure SQL サーバーを作成できます。
- B. Subscnption1 のどこにも Azure SQL サーバーを作成することはできません。
- C. Azure SQL サーバーは ContosoRG1 にのみ作成できます。
- D. ContosoRG1 のみで Azure SQL Server を作成することはできません。

正解: (正解を表示します)

質問: 127

注: この質問は、同じシナリオを提示する一連の質問の一部です。一連の質問にはそれぞれ、定められた目標を満たす独自の解決策が含まれています。質問セットによっては、正しい解決策が複数ある場合もあれば、正しい解決策がない場合もあります。

このセクションの質問に回答した後は、その質問に戻ることはできません。そのため、これらの質問はレビュー画面に表示されません。

image1 という名前のイメージを含む Registry1 という名前の Azure コンテナ レジストリがあります。

image1 を使用してコンテナ インスタンスをデプロイしようとする、エラー メッセージが表示されます。

image1 を使用してコンテナ インスタンスをデプロイする必要があります。

解決策: Registry1 の管理者ユーザーを有効に設定します。

これは目標を満たしていますか?

A. はい

B. いいえ

正解: [\(正解を表示します\)](#)

When deploying an Azure Container Instance (ACI) using an image from an Azure Container Registry (ACR), the deployment must authenticate to the registry to pull the image. By default, ACR is secured, and you cannot pull images anonymously unless they are public.

According to the Azure Administrator study guide, enabling the Admin user option in ACR generates two access credentials (username and password) that can be used for authentication when pulling container images. When the Admin user setting is disabled, you cannot authenticate directly using those credentials, and automated deployments may fail with authentication errors.

By enabling the Admin user setting, Azure creates a service-level user account for the registry with permissions to push and pull images. These credentials can then be referenced securely in your container instance configuration - either through the Azure portal, ARM template, or CLI parameters such as -- registry-username and --registry-password.

This configuration allows ACI to authenticate successfully against the ACR and pull the required image (image1) for deployment.

Therefore, enabling the Admin user for the container registry meets the goal of allowing the deployment of the container instance using image1.

質問: 128

初期ドメイン名を持つ Azure Active Directory (Azure AD) テナントがあります。

サードパーティのレジストラに contoso.com というドメイン名を登録しています。

名前に @contoso.com というサフィックスが含まれる Azure AD ユーザーを作成できることを確認する必要があります。

順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なコマンドレットをコマンドレットの一覧から回答領域に移動し、正しい順序に並べます。

Actions

Configure company branding.

Add an Azure AD tenant.

Verify the domain.

Create an Azure DNS zone.

Add a custom domain name.

Add a record to the public contoso.com DNS zone.

Answer Area

➡

⬅

⬆

⬆

正解:

Actions

- Configure company branding.
- Add an Azure AD tenant.
- Verify the domain.
- Create an Azure DNS zone.
- Add a custom domain name.
- Add a record to the public contoso.com DNS zone.

Answer Area

- Add a custom domain name.
- Add a record to the public contoso.com DNS zone.
- Verify the domain.

Explanation:

Answer Area

Add a custom domain name.

Add a record to the public contoso.com DNS zone.


Verify the domain.



The process is simple:  
Add the custom domain name to your directory  
Add a DNS entry for the domain name at the domain name registrar  
Verify the custom domain name in Azure AD  
References: <https://docs.microsoft.com/en-us/azure/dns/dns-web-sites-custom-domain>

質問: 129

ASP1 という名前の Azure App Service プランがあります。  
ASP1 の CPU 使用率は次の図に示されています。



ドロップダウンメニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

注意: 正しい選択ごとに1ポイントが加算されます。



正解:



Explanation:

The average CPU percentage is calculated 24 times per day. This is because the exhibit shows the CPU percentage for ASP1 in a 24-hour period, with one data point for each hour. Therefore, the average CPU percentage is calculated once per hour, or 24 times per day1.

ASP1 must be scaled out to optimize CPU usage. This is because the exhibit shows that the CPU percentage for ASP1 is consistently above 80%, which indicates that the app service plan is under high load and needs more instances to handle the traffic. Scaling out means adding more instances to an app service plan, which can improve the performance and availability of the apps hosted on it2. Scaling up means changing the pricing tier of an app service plan, which can increase the resources available for each instance, but not necessarily reduce the CPU usage3.

質問: 130

storage1 という名前のストレージ アカウントを含む Azure サブスクリプションがあります。

contains という名前の BLOB コンテナを作成する予定です。

包含には顧客管理キー暗号化を使用する必要があります。

どのキーを使用すればよいですか？

- A. P-384カーブのみを使用するECキー
- B. P-521カーブのみを使用するECキー
- C. P-384曲線またはP-521曲線のみを使用するECキー
- D. キーサイズが4096のRSAキー
- E. キーサイズが2048、3072、または4096のみのRSAキータイプ

正解: (正解を表示します)

When configuring customer-managed keys (CMK) for Azure Storage encryption, Microsoft requires the use of Azure Key Vault-managed keys that meet specific cryptographic standards. According to the Azure Storage security and encryption documentation, only RSA keys are supported for customer-managed encryption keys. Elliptic Curve (EC) keys, regardless of curve type (P-384 or P-521), are not supported for Azure Storage CMK scenarios.

Azure Storage supports RSA keys with the following key sizes:

- \* 2048-bit
- \* 3072-bit
- \* 4096-bit

These key sizes align with Microsoft's encryption compliance and security baseline requirements. When a storage account is configured to use CMK, all supported services-including Blob containers-inherit encryption using the specified RSA key from Azure Key Vault.

Because the requirement is to encrypt the blob container contained using customer-managed keys, the only valid and supported choice is an RSA key with one of the supported key sizes listed above.

Final Answer: E. an RSA key type with a key size of 2048, 3072, or 4096 only

質問: 131

次の表に示すユーザーを含む contoso.onmicrosoft.com という名前の Microsoft Entra テナントがあります。

| Name  | Member of      | Role assigned      |
|-------|----------------|--------------------|
| User1 | Group1         | None               |
| User2 | Group2         | None               |
| User3 | Group1, Group2 | User Administrator |

パスワード リセットの図に示すように、contoso.onmicrosoftxom のパスワード リセットを有効にします。([パスワード リセット] タブをクリックします。)



認証方法の図に示すように、パスワード リセットの認証方法を構成します。  
([認証方法]タブをクリックします。)

Number of methods required to reset ⓘ

1 2

Methods available to users

☐ Mobile app notification

☐ Mobile app code

☐ Email

☒ Mobile phone

☐ Office phone

☒ Security questions

Number of questions required to register ⓘ

3 4 5

Number of questions required to reset ⓘ

3 4 5

Select security questions >

10 security questions selected

**i** These settings only apply to end users in your organization. Admins are always enabled for self-service password reset and are required to use two authentication methods to reset their password. [Click here to learn more about administrator password policies.](#)

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。  
注意: 正しい選択ごとに1ポイントが加算されます。

Answer Area

Statements

Yes

No

After User2 answers three security questions correctly, he can reset his password immediately.

☐☐

If User1 forgets her password, she can reset the password by using the mobile phone app.

☐☐

User3 can add security questions to the password reset process.

☐☐

正解:

Answer Area

Statements

Yes

No

After User2 answers three security questions correctly, he can reset his password immediately.

☐☒

If User1 forgets her password, she can reset the password by using the mobile phone app.

☐☒

User3 can add security questions to the password reset process.

☒☐

Explanation:

| Statements                                                                                     | Yes                              | No                               |
|------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| After User2 answers three security questions correctly, he can reset his password immediately. | <input type="radio"/>            | <input checked="" type="radio"/> |
| If User1 forgets her password, she can reset the password by using the mobile phone app.       | <input type="radio"/>            | <input checked="" type="radio"/> |
| User3 can add security questions to the password reset process.                                | <input checked="" type="radio"/> | <input type="radio"/>            |

質問: 132

ストレージ アカウントを含む Azure サブスクリプションがあります。

Window Server 2016 を実行する Server1 という名前のオンプレミス サーバーがあります。Server1 には 2 TB のデータがあります。

Azure Import/Export サービスを使用して、データをストレージ アカウントに転送する必要があります。

どの順序でアクションを実行すればよいですか? 回答するには、すべてのアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

注意：正解の選択肢の順序は複数あります。正解の選択肢のいずれかを選択しても、得点は付与されます。

Actions

Detach the external disks from Server1 and ship the disks to an Azure data center.

From the Azure portal, update the import job.

Attach an external disk to Server1 and then run waimportexport.exe.

From the Azure portal, create an import job.

Answer Area

正解:

Actions

Detach the external disks from Server1 and ship the disks to an Azure data center.

From the Azure portal, update the import job.

Attach an external disk to Server1 and then run waimportexport.exe.

From the Azure portal, create an import job.

Answer Area

Attach an external disk to Server1 and then run waimportexport.exe.

From the Azure portal, create an import job.

Detach the external disks from Server1 and ship the disks to an Azure data center.

From the Azure portal, update the import job.

Explanation:

Answer Area

1 Attach an external disk to Server1 and then run waimportexport.exe.

2 From the Azure portal, create an import job.

3 Detach the external disks from Server1 and ship the disks to an Azure data center.

4 From the Azure portal, update the import job.

At a high level, an import job involves the following steps:

Step 1: Attach an external disk to Server1 and then run waimportexport.exe Determine data to be imported, number of drives you need, destination blob location for your data in Azure storage. Use the WAImportExport tool to copy data to disk drives. Encrypt the disk drives with BitLocker.

Step 2: From the Azure portal, create an import job. Create an import job in your target storage account in Azure portal. Upload the drive journal files.

Step 3: Detach the external disks from Server1 and ship the disks to an Azure data center. Provide the return address and carrier account number for shipping the drives back to you. Ship the disk drives to the shipping address provided during job creation.

Step 4: From the Azure portal, update the import job Update the delivery tracking number in the import job details and submit the import job. The drives are received and processed at the Azure data center. The drives are shipped using your carrier account to the return address provided in the import job.

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-import-export-service>

質問: 133

Azureサブスクリプションには、storage1という名前のストレージアカウントが含まれています。storage1アカウントには、containet1という名前のコンテナが含まれています。

rule1 という名前の BLOB ライフサイクル ルールを作成します。

45 日間更新されなかった BLOB をコンテナから自動的に移動するように、rule1 を構成する必要があります。

クール アクセス ティアへ。

ルールをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正解ごとに 1 ポイントが加算されます。

Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToCool": {
              "daysAfterCreationGreaterThan": 45,
              "daysAfterLastAccessTimeGreaterThan": 45,
              "daysAfterModificationGreaterThan": 45
            }
          }
        },
        "filters": {
          "blobTypes": [
            "AppendBlob",
            "Blockblob",
            "Pageblob"
          ],
          "prefixMatch": [
            "container1"
          ]
        }
      }
    }
  ]
}
```

正解:

## Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToCool": {
              "daysAfterCreationGreaterThan": 45,
              "daysAfterLastAccessTimeGreaterThan": 45,
              "daysAfterModificationGreaterThan": 45
            }
          }
        },
        "filters": {
          "blobTypes": [
            "AppendBlob",
            "Blockblob",
            "Pageblob"
          ],
          "prefixMatch": [
            "container1"
          ]
        }
      }
    }
  ]
}
```

Explanation:



Azure Blob Storage Lifecycle Management enables you to automatically transition data to cooler storage tiers, archive data, or delete data based on specified conditions related to creation, last modification, or last access time. These policies are defined in a JSON-based rule structure under the rules property.

As per Microsoft Azure Administrator Study Guide and Official Documentation ("Manage the Azure Blob storage lifecycle"), each rule must define:

Actions: What operation to perform (e.g., tierToCool, tierToArchive, delete).

Filters: Which objects (blobs) the rule applies to, typically filtered by container name, prefix, or blob type.

Conditions: When to perform the action, defined by properties such as:

"daysAfterCreationGreaterThan" - Number of days since blob creation.

"daysAfterLastAccessTimeGreaterThan" - Number of days since last read operation (used when last access tracking is enabled).

"daysAfterModificationGreaterThan" - Number of days since last modification (used for update or write operations).

According to the scenario, the requirement is to move blobs that were not updated for 45 days. The "not updated" condition refers to the last modification date, not creation or access date. Therefore, the correct property to use is:

"daysAfterModificationGreaterThan" : 45

The blob type must be specified under "filters". Azure Lifecycle Management currently supports lifecycle actions for:

BlockBlob - the primary type for uploaded files.

AppendBlob - supported, but typically used for log data.

PageBlob - not supported for lifecycle tiering (used for virtual machine disks).

Per Microsoft Docs - "Lifecycle management policy definition", lifecycle tiering actions (tierToCool or tierToArchive) only apply to BlockBlob and AppendBlob, with BlockBlob being the standard choice for data that benefits from tier transitions between Hot, Cool, and Archive tiers.

Hence, the correct configuration is:

"tierToCool" : {

"daysAfterModificationGreaterThan" : 45

```
}
" blobTypes " : [
" BlockBlob "
]
```

This policy automatically transitions block blobs in container1 to the Cool access tier when they have not been modified for 45 days, optimizing storage cost while retaining availability.

質問: 134

cont2 に対して計画された変更を実装します。

cont2 に作成できる追加アクセス ポリシーの最大数はいくつですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:

Stored access policies: # 2

Immutable blob storage policies: # 1

From the scenario in the earlier case study (A. Datum Corporation):

- \* The planned change requires creating a new container named cont2 in storage1, with the following:
- \* Three stored access policies (Stored1, Stored2, Stored3)
- \* A legal hold for immutable blob storage

Also, recall from the table:

- \* storage1 has Hierarchical namespace = Yes (Data Lake Storage Gen2 enabled).
- \* storage2 has Hierarchical namespace = No.

The question asks: "What is the maximum number of additional access policies you can create for cont2?" According to Microsoft Azure Storage documentation:

- \* A blob container can have a maximum of five stored access policies.
- \* These policies allow shared access signatures (SAS) to be managed centrally, letting you define start times, expiry times, and permissions at the container level.

Since three stored access policies (Stored1, Stored2, Stored3) already exist, the maximum additional policies that can still be created is:

# 5 (maximum allowed) - 3 (already existing) = 2

For immutable blob storage policies:

- \* A container can have one active immutability policy at a time.
- \* This can be either a time-based retention policy or a legal hold policy.
- \* Since cont2 already has a legal hold applied, no additional immutable policy can coexist with it, but it can have one defined policy type (the legal hold).

Therefore:

- \* Stored access policies: 2 additional policies can still be created.
- \* Immutable blob storage policies: 1 policy (the existing legal hold).

This aligns exactly with Microsoft Learn: Azure Storage Blob Service limits and immutable storage documentation:

"A container may have up to five stored access policies. Immutable storage supports either a time-based retention policy or legal hold policy per container." Final Verified Answer:

# Stored access policies: 2

# Immutable blob storage policies: 1

質問: 135

注: この質問は、同じシナリオを提示する一連の質問の一部です。一連の質問にはそれぞれ、定められた目標を満たす可能性のある独自の解決策が含まれています。質問セットによっては、正しい解決策が複数ある場合もあれば、正しい解決策がない場合もあります。

このセクションで質問に答えた後は、そのセクションに戻ることはできません。そのため、これらの質問はレビュー画面に表示されません。

image1 という名前のイメージを含む Registry1 という名前の Azure コンテナ レジストリがあります。

image1 を使用してコンテナ インスタンスをデプロイしようとすると、エラー メッセージが表示されます。

image1 を使用してコンテナ インスタンスをデプロイする必要があります。

解決策: Registry1 に専用データ エンドポイントを使用するを選択します。

これは目標を満たしていますか？

A. はい

B. いいえ

正解: ([正解を表示します](#))

In Microsoft Azure, the Azure Container Registry (ACR) is a managed service that allows you to store and manage container images for all types of container deployments. When deploying an Azure Container Instance (ACI) from an ACR image, the deployment may fail if network configuration, authentication, or permissions are not correctly set.

The option "Use dedicated data endpoint" in ACR is designed to isolate network traffic for data operations (such as image push/pull) from the registry's control plane operations. However, enabling or disabling this feature does not affect authentication or deployment permissions to an Azure Container Instance.

According to the Azure Administrator Study Guide (Microsoft Official Documentation):

"To deploy a container instance from an Azure container registry, the registry must be accessible either publicly with proper authentication (admin user or service principal with AcrPull permission) or privately using a Virtual Network with Private Link (Premium tier). If you receive authentication or access errors, the solution is to verify credentials or network accessibility, not to enable a dedicated data endpoint." In this case, the root cause of the deployment error is most likely related to image access authentication or tier limitations, not the use of dedicated data endpoints. Therefore, selecting Use dedicated data endpoint will not resolve the deployment failure.

The verified solution as per Microsoft Learn and AZ-104 exam content is to either:

Enable the admin user or

Assign a managed identity or service principal with the AcrPull role to the container instance.

Hence, the proposed solution does not meet the goal.

質問: 136

storage1という名前のAzureストレージアカウントがあります。

AzCopyを使用してデータをstorage1にコピーする予定です。

データをコピーできるstorage1のストレージサービスを識別する必要があります。

何を識別すべきですか？

A. blob、ファイル、テーブル、およびキュー

B. BLOBとファイルのみ

C. ファイルとテーブルのみ

D. ファイルのみ

E. BLOB、テーブル、およびキューのみ

正解: B (コメントを發表する)

AzCopy is a command-line utility provided by Microsoft Azure for transferring data to and from Azure Storage accounts with high performance and reliability. It is particularly designed to support bulk data migration and synchronization for Azure storage services.

According to the Microsoft Azure Administrator documentation (AzCopy v10 and later), the tool supports copying data to and from only two Azure storage services:

\* Azure Blob storage - including block blobs, append blobs, and page blobs.

\* Azure Files (File shares) - including directories and individual files.

The Table and Queue storage services are not supported by AzCopy, as these services handle structured and message-based data rather than binary or file data.

From the official Azure Administrator Study Guide and Microsoft Docs:

"AzCopy is a command-line tool that moves data into and out of Azure Storage accounts. You can use AzCopy with Blob storage or Azure Files. AzCopy doesn't support copying data to or from Table storage or Queue storage." (Source: Microsoft Learn - "Transfer data with AzCopy and Blob storage" and "Transfer data with AzCopy and Azure Files")

有効的な**AZ-104J**問題集はJPNTTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 137

contoso.com という名前の Azure Active Directory (Azure AD) テナントがあります。

fabrilcam.com と litwareinc.com という 2 つの外部パートナー組織があります。FabtAam.com は接続組織として構成されています。

アクセス パッケージの展示に示されているように、アクセス パッケージを作成します。(アクセス パッケージ ラボをクリックします。)

# New access package ...

\* Basics   Resource roles   \* Requests   Requestor information   \* Lifecycle   Review + Create

## Summary of access package configuration

### Basics

Name package1  
Description Guest users  
Catalog name General

### Resource roles

| Resource | Type           | Sub Type       | Role   |
|----------|----------------|----------------|--------|
| Group1   | Group and Team | Security Group | Member |

### Requests

Users who can request access All configured connected organizations  
Require approval No  
Enabled Yes

### Requestor information

#### Questions

| Question | Answer format | Multiple choice optio... | Required |
|----------|---------------|--------------------------|----------|
|----------|---------------|--------------------------|----------|

#### Attributes (Preview)

| Attribute type | Attribute | Default display string | Answer format | Multi |
|----------------|-----------|------------------------|---------------|-------|
|----------------|-----------|------------------------|---------------|-------|

### Lifecycle

Access package assignments expire After 365 days  
Require access reviews No

外部ユーザーのライフサイクル設定は、「ライフサイクル」の図に示すように構成します。（「ライフサイクル」タブをクリックします）

### Manage the lifecycle of external users

Select what happens when an external user, who was added to your directory through an access package request, loses their last assignment to any access package.

Block external user from signing in to this directory

Yes

No

Remove external user

Yes

No

Number of days before removing external user from this directory

30

### Delegate entitlement management

By default, only Global Administrators and User Administrators can create and manage catalogs, and can manage all catalogs. Users added to entitlement management as Catalog creators can also create catalogs and will become the owner of any catalogs they create.

Catalog creators ⓘ

0 selected

Add catalog creators

次の各文について、正しい場合は「はい」を選択し、そうでない場合は「いいえ」を選択します。注: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Statements

Litwareinc.com users can be assigned to package1.

After 365 days, fabrikam.com users will be removed from Group1.

After 395 days, fabrikam.com users will be removed from the contoso.com tenant.

Yes

No

正解:

Answer Area

Statements

Litwareinc.com users can be assigned to package1.

After 365 days, fabrikam.com users will be removed from Group1.

After 395 days, fabrikam.com users will be removed from the contoso.com tenant.

Yes

No

Explanation:

Litwareinc.com users can be assigned to package1. = No

After 365 days, fabrikam.com users will be removed from Group1. = Yes

After 395 days, fabrikam.com users will be removed from the contoso.com tenant = No

Litwareinc.com users cannot be assigned to package1 because they are not a connected organization in the contoso.com tenant. Only users from connected organizations can request access packages that are configured for external users1

Fabrikam.com users will be removed from Group1 after 365 days because the access package has an expiration policy of 365 days for external users. This means that the access assignments for external users will end after 365 days, unless they are renewed or extended2

Fabrikam.com users will not be removed from the contoso.com tenant after 395 days because the external user lifecycle settings have a deletion policy of 30 days after blocking. This means that external users will be blocked from signing in after 365 days of inactivity, and then deleted after another 30 days. Therefore, the total time before deletion is 395 days of inactivity, not 395 days from the date of assignment3

質問: 138

Azureサブスクリプションがあり、その中には米国東部2リージョンにVNETという仮想ネットワークが含まれています。VM1-NIというネットワークインターフェイスがVNET1に接続されています。

次の Azure Resource Manager テンプレートを正常にデプロイしました。

```

{
  "apiVersion": "2017-03-30",
  "type": "Microsoft.Compute/virtualMachines",
  "name": "VM1",
  "zones": "1",
  "location": "EastUS2",
  "dependsOn": [
    "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
  ],
  "properties": {
    "hardwareProfile": {
      "vmSize": "Standard_A2_v2"
    },
    "osProfile": {
      "computerName": "VM1",
      "adminUsername": "AzureAdmin",
      "adminPassword": "[parameters('adminPassword')]"
    },
    "storageProfile": {
      "imageReference": "[variables('image')]",
      "osDisk": {
        "createOption": "FromImage"
      }
    },
    "networkProfile": {
      "networkInterfaces": [
        {
          "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
        }
      ]
    }
  }
},
{
  "apiVersion": "2017-03-30",
  "type": "Microsoft.Compute/virtualMachines",
  "name": "VM2",
  "zones": "2",
  "location": "EastUS2",
  "dependsOn": [
    "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
  ],
  "storageProfile": {
    "imageReference": "[variables('image')]",
    "osDisk": {
      "createOption": "FromImage"
    }
  },
  "networkProfile": {
    "networkInterfaces": [
      {
        "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
      }
    ]
  }
}
]
}

```

Answer Area

Yes No

VM1 and VM2 can connect to VNET1.

☐ ☐

If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.

☐ ☐

If the East US 2 region becomes unavailable, VM1 or VM2 will be available.

☐ ☐

正解:

Answer Area



Microsoft

Yes

No

VM1 and VM2 can connect to VNET1.

☒☐

If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.

☒☐

If the East US 2 region becomes unavailable, VM1 or VM2 will be available.

☐☒

Explanation:

Answer Area

VM1 and VM2 can connect to VNET1.

Yes

No

If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.

Yes

No

If the East US 2 region becomes unavailable, VM1 or VM2 will be available.

Yes

No

" A resource can only be created in a virtual network that exists in the same region and subscription as the resource. " <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-vnet-plan-design-arm#regions>

質問: 139

VNET1 という名前の仮想ネットワークを含む Azure サブスクリプションがあります。VNET1 には、次の表に示すサブネットが含まれています。

| Name    | Connected virtual machines |
|---------|----------------------------|
| Subnet1 | VM1, VM2                   |
| Subnet2 | VM3, VM4                   |
| Subnet3 | VM5, VM6                   |

各仮想マシンは静的 IP アドレスを使用します。

次の要件を満たすには、ネットワーク セキュリティ グループ (NSG) を作成する必要があります。

- インターネットから VM3、VM4、VM5、VM6 への Web 要求を許可します。
- VM1 と VM2 間のすべての接続を許可します。
- VM1 へのリモート デスクトップ接続を許可します。
- VNET1 へのその他のすべてのネットワーク トラフィックを防止します。

作成する必要がある NSG の最小数はいくつですか？

- A. 12
- B. 4
- C. 1
- D. 3

正解: C ([コメントを发表する](#))

質問: 140

次の表に示す Azure コンテナ レジストリを持つ Azure サブスクリプションがあります。

| Name     | Service tier |
|----------|--------------|
| ContReg1 | Premium      |
| ContReg2 | Standard     |
| ContReg3 | Basic        |

ACR タスクを使用し、エンドポイント接続を構成する予定です。

Answer Area

ACR Tasks:

ContReg1 only  
ContReg1 and ContReg2 only  
ContReg1, ContReg2, and ContReg3

Private endpoints:

ContReg1 only  
ContReg1 and ContReg2 only  
ContReg1, ContReg2, and ContReg3

正解:

Answer Area

ACR Tasks:

ContReg1 only  
ContReg1 and ContReg2 only  
ContReg1, ContReg2, and ContReg3

Private endpoints:

ContReg1 only  
ContReg1 and ContReg2 only  
ContReg1, ContReg2, and ContReg3

Explanation:

# ACR Tasks: ContReg1 and ContReg2 only

# Private endpoints: ContReg1 only

Azure Container Registry (ACR) provides multiple service tiers (SKUs)-Basic, Standard, and Premium- each offering different capabilities for scalability, performance, and features. These tiers determine what functionalities are available, including ACR Tasks and Private Link (Private Endpoints) integration.

Here's how each tier differs according to Microsoft Azure Administrator documentation (AZ-104 Exam Study Guide and Azure Docs):

Feature

Basic

Standard

Premium

ACR Tasks (Build automation, image building, and updating)

# Supported

# Supported

# Supported

Private Link / Private Endpoints (Private access via Azure backbone)

# Not Supported

# Not Supported

# Supported

Geo-replication

# Not Supported

# Not Supported

# Supported

Content trust and RBAC integration

# Supported

# Supported

# Supported

Analysis:

ACR Tasks:

ACR Tasks allow you to automate image builds and updates when base images or application code changes.

Supported in Standard and Premium tiers (as per Microsoft Docs: "ACR Tasks are available in Standard and Premium service tiers.").

Therefore, ContReg1 (Premium) and ContReg2 (Standard) support ACR Tasks.

ContReg3 (Basic) does not support advanced tasks automation at scale.

# Answer: ContReg1 and ContReg2 only

Private Endpoints:

Private Link (Private Endpoints) enables private network connectivity to the ACR from a virtual network.

According to Azure documentation:

"Private Link for Azure Container Registry is supported only in the Premium service tier." This ensures that access to the registry occurs via Azure's backbone, preventing exposure to the public internet.

# Answer: ContReg1 only (Premium)

Verified Microsoft Documentation Extract:

From Microsoft Learn: Azure Container Registry service tiers overview:

"The Premium tier provides advanced capabilities such as geo-replication, content trust, customer-managed keys, and Private Link support."

"ACR Tasks is available in both Standard and Premium tiers for continuous integration workflows and automated image building."

質問: 141

注: この質問は、同じシナリオを示す一連の質問の一部です。このシリーズの各質問には、指定された目標を達成できる可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策が含まれる場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答すると、その質問に戻ることはできません。そのため、これらの質問はレビュー画面には表示されません。

10 個の仮想ネットワークを含む Azure サブスクリプションがあります。仮想ネットワークは別のリソース グループでホストされます。

別の管理者は、サブスクリプション内に複数のネットワーク セキュリティ グループ (NSG) を作成することを計画しています。

NSG の作成時に、仮想ネットワーク間の TCP ポート 8080 が自動的にブロックされるようにする必要があります。

解決策: カスタム ポリシー定義を構成し、そのポリシーをサブスクリプションに割り当てます。

これは目標を達成していますか？

A. はい

B. いいえ

正解: [\(正解を表示します\)](#)

A custom policy definition is a way to define your own rules for using Azure resources. You can use custom policies to enforce compliance, security, cost management, or organization-specific requirements. However, a custom policy definition alone is not enough to meet the goal of automatically blocking TCP port 8080 between the virtual networks. You also need to create a policy assignment that applies the custom policy definition to the scope of the subscription. A policy assignment is the link between a policy definition and an Azure resource. Without a policy assignment, the custom policy definition will not take effect. Therefore, the solution does not meet the goal.

References:

Tutorial: Create a custom policy definition

Create and manage policies to enforce compliance

質問: 142

次の表に示すグループを含む、adatum.com という名前の Azure AD テナントがあります。

| Name   | Member of |
|--------|-----------|
| Group1 | None      |
| Group2 | Group1    |
| Group3 | Group2    |

Adatum.com には、次の表に示すユーザーが含まれています。

| Name  | Member of |
|-------|-----------|
| User1 | Group1    |
| User2 | Group2    |
| User3 | Group3    |
| User4 | None      |

Azure AD Premium P2 ライセンスをグループ1 と User4 に割り当てます。  
Azure AD Premium P2 ライセンスはどのユーザーに割り当てられていますか？

- A. ユーザー4のみ
- B. ユーザー1とユーザー4のみ
- C. User1、User2、User4のみ
- D. ユーザー1、ユーザー2、ユーザー3、ユーザー4

正解: B ([コメントを发表する](#))

According to the Microsoft documentation, when you assign a license to a group, all members of that group are automatically assigned the license. However, if a user is already assigned the same license directly or through another group, the license is not duplicated.

In your scenario, you assigned the Azure AD Premium P2 license to Group1 and User4. This means that all members of Group1, which are User1 and User2, will also get the license. User4 will get the license directly.

User3 will not get the license because they are not a member of Group1 or assigned the license directly.


Therefore, the users who are assigned the Azure AD Premium P2 license are User1, User2, and User4 only.

質問: 143

次の表に示すリソースを含む Azure サブスクリプションがあります。

| Name   | Type                    | Resource group | Region  |
|--------|-------------------------|----------------|---------|
| Vault1 | Recovery services vault | RG1            | East US |
| VM1    | Virtual machine         | RG1            | East US |
| VM2    | Virtual machine         | RG1            | West US |

すべての仮想マシンは Windows Server を実行します。  
VM1では、次の図に示すように、Folder!という名前のフォルダをバックアップします。



バックアップを別の仮想マシンに復元する予定です。

バックアップを VM2 に復元する必要があります。

まず何をすべきでしょうか？

- A. VM1からMicrosoft Azure Recovery Servicesエージェントをインストールします
- B. VM1 から Windows Server バックアップ機能をインストールします。
- C. VM2 から、Windows Server バックアップ機能をインストールします。
- D. VM2 から、Microsoft Azure Recovery Services エージェントをインストールします。

正解: D ([コメントを发表する](#))

The screenshot shows the Schedule Backup Wizard (Files and Folders), which is the Microsoft Azure Recovery Services (MARS) agent experience. This is file/folder backup to a Recovery Services vault, not an "Azure VM backup" policy-based backup. In Microsoft's Azure Backup guidance, the MARS agent protects files and folders on a Windows machine by uploading encrypted backup data to the Recovery Services vault.

The encryption is controlled by a customer-managed passphrase, and Azure does not store that passphrase.

Because MARS backups are agent-based, restore is not limited to the original VM or to the VM's region.

Microsoft documentation for MARS restore describes restoring to the original location or to an alternate location, including another server, as long as the target machine is registered to the same vault and you provide the same passphrase used to encrypt the backups. The presence of VM2 in West US does not block restore for MARS-protected files/folders; region coupling is a constraint for Azure VM backup (vault and VM must be in the same region), not for MARS agent file/folder backup.

So, after redeploying or on another VM, you install and register the MARS agent to Vault1, then perform a restore of Folder1 using the correct passphrase-VM2 is a valid restore target.

質問: 144

次の表に示す仮想マシンを含む Azure サブスクリプションがあります。

| Name | Operating system    | Connects to |
|------|---------------------|-------------|
| VM1  | Windows Server 2019 | Subnet1     |
| VM2  | Windows Server 2019 | Subnet2     |

VM1 と VM2 はパブリック IP アドレスを使用します。VM1 と VM2 上の Windows Server 2019 から、受信リモートデスクトップ接続を許可します。

Subnet1 と Subnet2 は、VNET1 という名前の仮想ネットワーク内にあります。

サブスクリプションには、NSG1 と NSG2 という 2 つのネットワーク セキュリティ グループ (NSG) が含まれています。NSG1 はデフォルトのルールのみを使用します。

NSG2 はデフォルトのルールと次のカスタム受信ルールを使用します。

- \* 優先度: 100
- \* 名前: ルール1
- \* ポート: 3389
- \* プロトコル: TCP
- \* 出典: 任意
- \* 目的地: 任意
- \* アクション: 許可

NSG1 はサブネットに関連付けられています。NSG2 は VM2 のネットワーク インターフェイスに関連付けられています。

以下の各文について、正しい場合は 「はい」を選択してください。そうでない場合は 「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

| Answer Area | Statements                                                         | Yes                   | No                    |
|-------------|--------------------------------------------------------------------|-----------------------|-----------------------|
|             | From the internet, you can connect to VM1 by using Remote Desktop. | <input type="radio"/> | <input type="radio"/> |
|             | From the internet, you can connect to VM2 by using Remote Desktop. | <input type="radio"/> | <input type="radio"/> |
|             | From VM1, you can connect to VM2 by using Remote Desktop.          | <input type="radio"/> | <input type="radio"/> |

正解:

Answer Area

Microsoft

Statements

Yes

No

From the internet, you can connect to VM1 by using Remote Desktop.

☐

☒

From the internet, you can connect to VM2 by using Remote Desktop.

☒

☐

From VM1, you can connect to VM2 by using Remote Desktop.

☒

☐

Explanation:

Answer Area

Microsoft

Statements

Yes

No

From the internet, you can connect to VM1 by using Remote Desktop.

☐

☒

From the internet, you can connect to VM2 by using Remote Desktop.

☒

☐

From VM1, you can connect to VM2 by using Remote Desktop.

☒

☐

In Azure, Network Security Groups (NSGs) control inbound and outbound traffic to network interfaces (NICs), subnets, and virtual machines (VMs) using rules based on priority and direction.

According to the Microsoft Azure Administrator Guide and Azure Networking documentation, the following principles apply:

Default NSG rules:

By default, an NSG denies all inbound traffic from the Internet except traffic originating from the same Virtual Network (VNet).

NSG allows all outbound traffic to the Internet.

Default inbound rules include:

Allow VNet inbound (priority 65000)

Allow Azure Load Balancer inbound (priority 65001)

Deny all inbound (priority 65500)

NSG associations:

NSGs can be associated either with a subnet or an individual network interface (NIC).

When both are applied, the NIC-level NSG takes precedence for inbound/outbound traffic.

If no explicit allow rule exists, default deny applies.

Analysis of Each Statement

1## From the internet, connect to VM1 by using RDP:

VM1's subnet (Subnet1) is associated with NSG1, which has only the default rules.

The default rules deny all inbound traffic from the Internet, including port 3389 (RDP).

# Therefore, RDP from the Internet is blocked.

The Answer: # No

2## From the internet, connect to VM2 by using RDP:  
VM2's NIC is associated with NSG2, which includes a custom allow rule (priority 100) permitting TCP traffic on port 3389 from any source to any destination.  
This rule overrides the default deny rule.  
# Thus, RDP from the Internet is allowed.  
The Answer: # Yes  
3## From VM1, connect to VM2 by using RDP:  
Both VMs reside in the same VNet (VNET1) but different subnets (Subnet1 and Subnet2).  
The default NSG rule "Allow VNet Inbound" allows traffic between subnets within the same virtual network.  
# Therefore, VM1 can connect to VM2 via RDP (port 3389).  
The Answer: # Yes

質問: 145

Azureサブスクリプションに、storage1というストレージアカウントが含まれています。storage1アカウントには、画像を保存するBLOBが含まれています。  
ストレージ 1 へのクライアント アクセスは、共有アクセス署名 (SAS) を使用して許可されます。  
ユーザーが 7 日間の期間を超える SAS を生成するときに、警告メッセージが表示されるようにする必要があります。  
ストレージ1はどうすればいいでしょうか？  
A. ライフサイクル管理ルールを追加します。  
B. 共有アクセス署名 (SAS) の有効期限間隔の推奨上限を許可するを [有効] に設定します。  
C. アラート ルールを構成します。  
D. 読み取り専用ロックを有効にします。

正解: [\(正解を表示します\)](#)  
Azure Storage provides a built-in security setting that warns users when creating a SAS token with an expiry exceeding seven days.  
When Allow recommended upper limit for SAS expiry interval is enabled:  
\* Users receive a warning message  
\* SAS creation is not blocked  
\* Enforces Microsoft's security best practice  
Lifecycle rules manage blob data, not access. Alerts cannot intercept SAS creation. Resource locks do not affect SAS behavior.  
Microsoft documentation states:  
"Enabling the recommended SAS expiry interval warns users when the expiry exceeds seven days."

質問: 146

storage1 というストレージアカウントを含む Azure サブスクリプションがあります。このサブスクリプションは、オンプレミスの Active Directory ドメインと同期する contoso.com という Azure Active Directory (Azure AD) テナントにリンクされています。  
ドメインには、次の表に示すセキュリティ プリンシパルが含まれています。

| Name      | Type     |
|-----------|----------|
| User1     | User     |
| Computer1 | Computer |

Azure AD で、User2 という名前のユーザーを作成します。  
storage1 アカウントには、share1 という名前のファイル共有が含まれており、次の構成になっています。

```
"kind": "StorageV2",
"properties": {
  "azureFilesIdentityBasedAuthentication": {
    "directoryServiceOptions": "AD",
    "activeDirectoryProperties": {
      "domainName": "Contoso.com",
      "netBiosDomainName": "Contoso.com",
      "forestName": "Contoso.com",
    }
  }
}
```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

| Statements                                                                                    | Yes                   | No                    |
|-----------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| You can assign the Storage File Data SMB Share Contributor role to User1 for share1.          | <input type="radio"/> | <input type="radio"/> |
| You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.           | <input type="radio"/> | <input type="radio"/> |
| You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1. | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                                                                                    | Yes                              | No                               |
|-----------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| You can assign the Storage File Data SMB Share Contributor role to User1 for share1.          | <input checked="" type="radio"/> | <input type="radio"/>            |
| You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.           | <input type="radio"/>            | <input checked="" type="radio"/> |
| You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1. | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

| Statements                                                                                    | Yes                              | No                               |
|-----------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| You can assign the Storage File Data SMB Share Contributor role to User1 for share1.          | <input checked="" type="radio"/> | <input type="radio"/>            |
| You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.           | <input type="radio"/>            | <input checked="" type="radio"/> |
| You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1. | <input checked="" type="radio"/> | <input type="radio"/>            |

Reference:

[https://docs.microsoft.com/en-us/azure/storage/files/storage-files-identity-ad-ds-assign-permissions?](https://docs.microsoft.com/en-us/azure/storage/files/storage-files-identity-ad-ds-assign-permissions?tabs=azure-portal)

tabs=azure-portal

Azure Files supports Azure Active Directory Domain Services (Azure AD DS) authentication for SMB access, allowing you to use your existing Active Directory (AD) identities and groups to control access to Azure file shares. In this scenario, the configuration snippet shows that the storage account (storage1) is using Active Directory (AD) authentication through azureFilesIdentityBasedAuthentication with directoryServiceOptions: " AD " . According to Microsoft Azure documentation, the Storage File Data SMB Share built-in roles (Reader, Contributor, Elevated Contributor) can only be assigned to Azure AD users, groups, or service principals, not to computer objects.

Statement 1: # User1 is an on-premises AD user synchronized to Azure AD. Since Azure Files uses identity- based authentication through AD, User1 can be assigned roles like Storage File Data SMB Share Contributor to access share1.

Statement 2: # Computer1 is a computer account and cannot have Azure RBAC roles directly assigned in Azure. Therefore, you cannot assign SMB Share Reader to it.

Statement 3: # User2 is an Azure AD user (cloud-only identity). Azure RBAC roles for storage shares can also be assigned to Azure AD users. Hence, you can assign Storage File Data SMB Share Elevated Contributor to User2.

**質問: 147**

次の表に示す Azure 仮想マシンがあります。

| Name | IP address | Virtual network |
|------|------------|-----------------|
| VM1  | 10.0.0.4   | VNET1           |
| VM2  | 10.0.0.5   | VNET1           |

VNET1 はプライベート DNS ゾーンにリンクされており、次の表に示すレコードを含む contoso.com という名前が付けられています。

| Name  | Type  | TTL  | Value             | Auto registered |
|-------|-------|------|-------------------|-----------------|
| comp1 | TXT   | 3600 | 10.0.0.5          | False           |
| comp2 | A     | 3600 | 10.0.0.5          | False           |
| comp3 | CNAME | 3600 | comp1.contoso.com | False           |
| comp4 | PTR   | 3600 | 10.0.0.5          | False           |

VM1 から VM2 に ping を実行する必要があります。

VM2 に ping を実行するために使用できる DNS 名は何ですか。

### A. comp2 contoso.com のみ

**B. comp1.contoso.com、comp2.contoso.com、comp4.contoso.con のみ**

**C. com1.contoso.com と comp2.contoso.com のみ**

**D. comp1.contoso.com、comp2.contoso.com、comp3.contoso.com、comp4.contoso.com**

E. comp2.contoso.com と comp4.contoso.com のみ

正解: (正解を表示します)

**質問: 148**

LB1 という名前の Azure ロード バランサーがあります。

User1 という名前のユーザーに、次の図に示すロールを割り当てます。

**User assignments - LB1**

Assignments for the selected user, group, service principal, or managed identity at this scope or inherited to this scope.

Search by assignment name or description

---

**Answer Area**

User1 can [answer choice] LB1.

- delete
- create a NAT rule for assign access to other users for

User1 can [answer choice] the resource group.

- delete a virtual machine from modify the load balancing rules in deploy an Azure Kubernetes Service (AKS) cluster to

正解:

Answer Area

User1 can [answer choice] LB1.

User1 can [answer choice] the resource group.

Microsoft

Explanation:

User1 can [answer choice] LB1.

User1 can [answer choice] the resource group.

User Access Administrator can only assign access to other users  
<https://docs.microsoft.com/en-us/azure/role-based-access-control/rbac-and-directory-admin-roles> Virtual Machine Contributor can Manage VMs, which includes deleting VMs too.  
<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles#virtual-machine-contributor>  
<https://docs.microsoft.com/en-us/answers/questions/350635/can-virtual-machine-contributor-create-vm.html>

質問: 149

技術要件を満たすには、VM1 と VM2 のアラートを構成する必要があります。  
順番に実行する必要がある 3 つのアクションはどれですか。回答するには、すべてのアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions

Configure the Diagnostic settings.

Collect Windows performance counters from the Log Analytics agents.

Create an alert rule.


Create an Azure SQL database.

Create a Log Analytics workspace.

Answer Area

Microsoft

正解:



Explanation:

1## Create a Log Analytics workspace.

2## Collect Windows performance counters from the Log Analytics agents.

3## Create an alert rule.

The question states:

"You need to configure the alerts for VM1 and VM2 to meet the technical requirements." The technical requirement from the case study specifies:

"Trigger an alert if VM1 or VM2 has less than 20 GB of free space on volume C." This requirement involves monitoring performance metrics (disk space) and generating alerts based on those metrics. According to Microsoft Azure monitoring and management documentation, the process to configure such alerts involves using Azure Monitor and Log Analytics.

Step-by-Step Verified Solution:

Step 1: Create a Log Analytics workspace

A Log Analytics workspace is required to store and analyze logs and performance data collected from virtual machines.

Azure Monitor uses this workspace as the central repository for performance counters, events, and logs from connected agents.

From the Microsoft Documentation:

"Before you can collect and query data from virtual machines, you must create a Log Analytics workspace in your subscription." (Source: Azure Monitor and Log Analytics Guide) Step 2: Collect Windows performance counters from the Log Analytics agents After creating the workspace, you must connect VM1 and VM2 to the workspace and configure the Windows performance counters that you want to monitor.

In this case, you would collect the LogicalDisk(%) Free Space counter for C: drive.

Azure Monitor agent or Log Analytics agent collects these metrics and sends them to the workspace for analysis.

"To monitor system performance, configure the agent to collect performance counters such as available memory or free disk space." (Source: Azure Monitor Performance Counters Documentation) Step 3: Create an alert rule

Once performance data is being collected, you can create an alert rule in Azure Monitor based on a Kusto query or metric threshold.

You would define a condition such as:

LogicalDisk | where FreeSpaceMB < 20480

and configure it to trigger an alert when free space on volume C drops below 20 GB.

This alert can notify via email, action group, or automation runbook.

"Alerts in Azure Monitor proactively notify you when important conditions are found in your monitoring data. Alerts can trigger automated actions or notifications." (Source: Azure Monitor Alerts Overview) Incorrect Options (Eliminated):

Configure the Diagnostic settings:

Used for collecting activity logs or resource logs, not performance counters from VMs.

Create an Azure SQL database:

Not relevant to the scenario of monitoring disk space.

質問: 150

Standard レベルの Azure コンテナ レジストリがあります。

管理者ユーザー アカウントを有効にして、アカウントのパスワードを取得する必要があります。どの設定を使用する必要がありますか？

A. ネットワーキング

B. プロパティ

C. アイデンティティ

D. アクセスキー

正解: (正解を表示します)

In Azure Container Registry (ACR), the admin user account provides a simple authentication mechanism using a username and password. Enabling the admin account and retrieving its credentials is done exclusively through the Access keys blade.

According to Azure Container Registry documentation:

The admin user is enabled/disabled from Access keys

The passwords (primary and secondary) are displayed and regenerated there Other blades such as Networking, Properties, or Identity do not manage admin credentials Final Answer:

D). Access keys

質問: 151

VM1 という仮想マシンを含む Azure サブスクリプションがあります。VM1 は、24 時間利用可能な基幹業務アプリケーションをホストしています。VM1 には、1 つのネットワーク インターフェイスと 1 つのマネージド ディスクがあります。VM1 は D4s v3 サイズを使用しています。

VM1 に次の変更を加える予定です。

\* サイズをD8s v3に変更します。

\* 500 GB の管理ディスクを追加します。

\* Puppet Agent 拡張機能を追加します。

\* Desired State Configuration Management を有効にします。

どの変更によって VM1 のダウンタイムが発生しますか？

A. Puppet エージェント拡張機能を追加します。

B. Desired State Configuration Management を有効にします。

C. サイズをD8s v3に変更します。

D. 500 GB の管理ディスクを追加します。

正解: (正解を表示します)

有効的な**AZ-104J**問題集はJPNTTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 152

オンプレミスネットワークには、Share1という名前のSMB共有が含まれています。

次のリソースを含むAzureサブスクリプションがあります。

-webapp1という名前のWebアプリ

-VNET1という名前の仮想ネットワーク

webapp1がShare1に接続できることを確認する必要があります。  
何をデプロイする必要がありますか？

- A. Azureアプリケーションゲートウェイ
- B. Azure Active Directory (Azure AD) アプリケーションプロキシ
- C. Azure仮想ネットワークゲートウェイ

正解: **C** ([コメントを發表する](#))

To enable a web app hosted in Azure App Service to connect securely to an on-premises SMB share (Share1), you must create hybrid network connectivity between your Azure environment and your on-premises network. According to the Microsoft Azure Administrator Study Guide and Microsoft Learn documentation, Azure Web Apps running in an App Service Plan cannot directly access on-premises file shares over the public internet for security reasons. You must extend your on-premises network to Azure through a Virtual Network (VNet) and then integrate the web app with that network.

The Virtual Network Gateway is the component that enables this hybrid connectivity. It establishes a Site-to-Site VPN or ExpressRoute connection between the Azure VNet and the on-premises network, allowing the web app (after VNet integration) to access internal resources such as SMB shares, SQL Servers, or file servers.

Once the VPN gateway is configured and the web app is integrated with the VNet (Regional VNet Integration), the web app can securely access Share1 over the private network channel.

Official Microsoft Documentation Extract (Summary):

"To access on-premises resources from Azure App Service, configure VNet Integration and establish a Site-to-Site VPN or ExpressRoute connection using a Virtual Network Gateway. This allows Azure resources to securely communicate with on-premises systems such as SMB file shares or databases." (Source: Microsoft Learn - Connect an App Service app to an on-premises network using Azure VPN Gateway.)

質問: **153**

User1 という名前のユーザーを含む Azure サブスクリプションがあります。

ユーザー1が仮想マシンをデプロイし、仮想ネットワークを管理できるようにする必要があります。このソリューションでは、最小権限の原則を適用する必要があります。

User1 に割り当てる必要があるロールベースのアクセス制御 (RBAC) ロールはどれですか？

- A. 仮想マシン コントリビューター
- B. 仮想マシン管理者ログイン
- C. 貢献者
- D. 所有者

正解: **C** ([コメントを發表する](#))

質問: **154**

VM1という名前のAzure仮想マシンを含むAzureサブスクリプションがあります。VM1は、複数のアクティブなインスタンスをサポートしないApp1という名前の財務報告アプリを実行します。

毎月末に、App1の実行時にVM1のCPU使用率がピークになります。

毎月末にVM1のプロセッサパフォーマンスを向上させるには、スケジュールされたRunbookを作成する必要があります。

Runbookにはどのようなタスクを含める必要がありますか？

- A. Azure Performance DiagnosticsエージェントをVM1に追加します。
- B. VM1のVMサイズプロパティを変更します。
- C. VM1をスケールセットに追加します。
- D. サブスクリプションのvCPU割り当てを増やします。
- E. 望ましい状態の構成 (DSC) 拡張機能をVM1に追加します。

正解: ([正解を表示します](#))

To create a scheduled runbook to increase the processor performance of VM1 at the end of each month, you need to modify the VM size property of VM1. This will allow you to scale up the VM to a larger size that has more CPU cores and memory. You can use Azure Automation to create a PowerShell runbook that changes the VM size using the Set-AzVM cmdlet. You can then schedule the runbook to run at the end of each month using the Azure portal or Azure PowerShell. For more information, see [How to resize a virtual machine in Azure using Azure Automation](#)1.

質問: 155

10台の仮想マシン、Vault 1というキーコンテナ、NSG1というネットワークセキュリティグループ (NSG)を含むAzureサブスクリプションがあります。すべてのリソースは、米国東部のAzureリージョンにデプロイされています。

仮想マシンは NSG1 を使用して保護されています。NSG1 は、インターネットへのすべての送信トラフィックをブロックするように構成されています。

仮想マシンがVault1にアクセスできるようにする必要があります。ソリューションでは、最小権限の原則を適用し、管理作業を最小限に抑える必要があります。

NSG1 の送信セキュリティ規則の宛先として何を構成する必要がありますか？

- A. サービスタグ
- B. アプリケーションセキュリティグループ
- C. IPアドレス範囲

正解: (正解を表示します)

In Azure, Network Security Groups (NSGs) control inbound and outbound traffic to network interfaces, subnets, and virtual machines. When NSGs block all outbound traffic, VMs cannot reach Azure services such as Azure Key Vault, Storage, or Azure SQL Database, unless exceptions are explicitly configured.

To allow virtual machines to access Vault1 (Azure Key Vault), the correct and most efficient configuration is to use an Azure service tag as the destination in the outbound rule.

- \* Service tags are predefined identifiers that represent a group of IP address prefixes for specific Azure services.
- \* Microsoft automatically manages the address prefixes included in each tag, reducing administrative overhead and maintenance.

The service tag for Azure Key Vault is:

AzureKeyVault

You can create an outbound rule in NSG1 with:

- \* Destination: Service tag # AzureKeyVault
- \* Port: 443 (HTTPS)
- \* Action: Allow

This allows only secure traffic to Azure Key Vault, satisfying the principle of least privilege (since it limits access to only the necessary service) while also minimizing administrative effort (since the service tag updates automatically as Azure changes).

Using an application security group or IP address range would require additional configuration and maintenance, making them less efficient.

# Final Verified Answer: A. a service tag

質問: 156

App1のソリューションを推奨する必要があります。ソリューションは技術要件を満たす必要があります。推奨内容にはどのような点を含めるべきですか？回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Number of virtual networks:

▼

1

2

3

Microsoft

Number of subnets:

▼

1

2

3

正解:

Number of virtual networks:

▼

1

2

3

Number of subnets:

▼

1

2

3



Explanation:

Number of virtual networks:

▼

1

2

3

Number of subnets:

▼

1

2

3

In this scenario, App1 consists of three distinct tiers - web front end, processing middle tier, and SQL database - each containing five virtual machines. The technical requirement specifies that the company must minimize the number of open ports between the App1 tiers, move all tiers of App1 to Azure, and ensure that all VMs are protected by backups.

According to Microsoft Azure architecture best practices for multi-tier applications (from Azure Architecture Center and the Azure Administrator curriculum), the optimal design involves:

- \* Deploying all tiers of App1 into a single virtual network (VNet).
- \* This allows all components of the application to communicate securely using private IP addresses.
- \* Keeping all tiers within a single VNet simplifies management, security, and monitoring while supporting Network Security Groups (NSGs) for inter-tier traffic control.
- \* Microsoft Documentation Extract:

"Use a single virtual network to host multi-tier applications. Divide the virtual network into multiple subnets, each representing a tier, and use network security groups (NSGs) to control traffic flow between tiers." (Source: Microsoft Learn - Design and implement virtual networks in Azure)

- \* Creating separate subnets for each application tier (3 total).
- \* Subnet 1: Web tier (internet-facing, HTTPS traffic)
- \* Subnet 2: Application/Processing tier (internal communication only)
- \* Subnet 3: Database tier (private, no internet access)
- \* Using NSGs, administrators can explicitly allow or deny traffic between subnets, thus minimizing open ports between tiers and meeting the security requirement.
- \* Microsoft Documentation Extract:

"Subnets provide isolation and segmentation within a virtual network. Each tier of an application should be deployed in its own subnet to apply network security policies and control exposure." (Source: Microsoft Learn - Azure virtual network design best practices)

- \* Backups and Storage Requirements:
- \* All VMs can use Azure Backup integrated with Recovery Services Vaults, which supports VM- level backup in a single VNet environment.
- \* The blueprint files are stored in Azure Blob Storage with the archive tier, ensuring compliance with the storage and access control requirements.

By using one virtual network and three subnets, Contoso ensures efficient management, minimized administrative overhead, secure isolation of application tiers, and full compliance with Azure governance and security recommendations.

# Final Verified Answer:

- \* Number of virtual networks: 1
- \* Number of subnets: 3

質問: 157

RG1 という名前のリソース グループを含む Azure サブスクリプションがあります。

storage1 という名前のストレージ アカウントを作成する予定です。

File1 という名前の Bicep ファイルがあります。

File1 を変更して、storage1 から RG1 へのデプロイメントを自動化できるようにする必要があります。

どのプロパティを変更する必要がありますか？

- A. kind
- B. location
- C. scope
- D. sku

正解: (正解を表示します)

質問: 158

Windows Server を実行する VM1 という名前の Azure 仮想マシンがあります。VM1 の Azure Resource Manager (ARM) テンプレートを生成します。ARM テンプレートから VM2 という名前の仮想マシンをデプロイする予定です。VM2 のデプロイ中に構成できる項目は何ですか？

- A. 管理者ユーザー名
- B. オペレーティングシステム
- C. 仮想マシンのサイズ
- D. リソースグループ

正解: (正解を表示します)

質問: 159

お客様はAzureサブスクリプションをご利用中です。このサブスクリプションには、Windows Server 2016を実行する仮想マシンが含まれており、以下の表に示すように構成されています。

| Name | Virtual network | DNS suffix configured in Windows Server |
|------|-----------------|-----------------------------------------|
| VM1  | VNET2           | Contoso.com                             |
| VM2  | VNET2           | None                                    |
| VM3  | VNET2           | Adatum.com                              |

adatum.com という名前のパブリック Azure DNS ゾーンと、contoso.com という名前のプライベート Azure DNS ゾーンを作成します。

次の図に示すように、contoso.com 用の仮想ネットワークリンクを作成します。

link1

contoso.com

Save

Discard

Delete

Access Control (IAM)

Tags

Link name

link1

Link state

Completed

Provisioning state

Succeeded

Virtual network details

Virtual network id

/subscriptions/8372f433-2d0d-4361-b5ef-5b188fed87d0/resourceGroups/RG2/provi...

Virtual network

VNET2

Configuration

Enable auto registration

以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。  
注：正解ごとに1ポイントが加算されます。

Answer Area

Statements

When VM1 starts, a record for VM1 is added to the contoso.com DNS zone.

When VM2 starts, a record for VM2 is added to the contoso.com DNS zone.

When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.

Yes

No

正解:

Answer Area

Statements

When VM1 starts, a record for VM1 is added to the contoso.com DNS zone.

When VM2 starts, a record for VM2 is added to the contoso.com DNS zone.

When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.

Yes

No

Explanation:

Answer Area

Microsoft

Statements

|                                                                         |                                      |                                     |
|-------------------------------------------------------------------------|--------------------------------------|-------------------------------------|
| When VM1 starts, a record for VM1 is added to the contoso.com DNS zone. | <input checked="" type="radio"/> Yes | <input type="radio"/> No            |
| When VM2 starts, a record for VM2 is added to the contoso.com DNS zone. | <input checked="" type="radio"/> Yes | <input type="radio"/> No            |
| When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.  | <input type="radio"/> Yes            | <input checked="" type="radio"/> No |

All three VMs are in VNET2. Auto registration is enabled for private Azure DNS zone named contoso.com, which is linked to VNET2. So, VM1, VM2 and VM3 will auto-register their host records to contoso.com. None of the VM will auto-register to the public Azure DNS zone named adatum.com. You cannot register private IPs on the internet (adatum.com) Box 1: Yes Auto registration is enabled for private Azure DNS zone named contoso.com.

Box 2: Yes  
Auto registration is enabled for private Azure DNS zone named contoso.com.

Box 3: No  
None of the VM will auto-register to the public Azure DNS zone named adatum.com Reference:  
<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-instances>  
<https://docs.microsoft.com/en-us/azure/dns/private-dns-autoregistration>  
<https://docs.microsoft.com/en-us/azure/dns/private-dns-virtual-network-links>

質問: 160  
Azure サブスクリプションをお持ちです。  
50 台の仮想マシンを VMware vSphere からサブスクリプションに移行する予定です。  
Recovery Services コンテナを作成します。  
次に何をすべきでしょうか？  
A. 拡張ネットワークを設定します。  
B. 復旧計画を作成します。  
C. Open Virtualization Application (OVA) テンプレートを vSphere にデプロイします。  
D. 仮想ネットワークを設定します。  
正解: (正解を表示します)

To migrate virtual machines from VMware vSphere to Azure, you need to use Azure Migrate, which is a service that helps you assess and migrate your on-premises workloads to Azure. Azure Migrate uses an appliance that you deploy as an Open Virtualization Application (OVA) template to vSphere. The appliance discovers the virtual machines and sends metadata and performance data to Azure Migrate. You can then use Azure Migrate to assess the readiness, cost, and sizing of the virtual machines for migration. You can also use Azure Migrate to replicate and migrate the virtual machines to Azure. References:  
\* About Azure Migrate  
\* Prepare VMware servers for assessment and migration to Azure with Azure Migrate Server Migration

質問: 161  
Template1とTemplate2という名前の2つのAzure Resource Manager(ARM)テンプレートがあります。

現在、Template1のデプロイが失敗するのは、Template1内の仮想マシン拡張機能が、Template2で参照されているリソースに依存しているためです。  
Template1とTemplate2は単一のトランザクションとして同時にデプロイされるようにし、どちらかのデプロイが失敗した場合は全体のデプロイも失敗するようにする必要があります。両方のテンプレートは独立しており、再利用可能でなければなりません。  
あなたはどうすべきでしょうか？

- A. Template1 で、仮想マシン拡張機能の provisionAfterExtensions プロパティを設定します。
- B. Template1 で、仮想マシン拡張機能の dependsOn プロパティを設定します。
- C. Template2 のリソースを Template1 に追加します。
- D. Template1 と Template2 をリンクされたテンプレートとして参照する親テンプレートを作成します。

正解: **D** ([コメントを發表する](#))

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

The correct approach is to create a parent ARM template that references Template1 and Template2 as linked templates . Microsoft's ARM guidance states that complex solutions can be broken into related templates and deployed together through a main template; linked templates remain separate files referenced from the main template, which preserves modularity and reusability.  
This also solves the cross-template dependency issue. Microsoft states that dependencies can be set between nested or linked templates, and if resources in one linked template must exist before resources in another, the second linked deployment can depend on the first. A simple dependsOn on the VM extension inside Template1 is insufficient because ARM dependencies are evaluated for resources deployed in the same template context; the documented dependency behavior is that Resource Manager deploys resources in dependency order and otherwise deploys them in parallel.  
Adding Template2's resources into Template1 would violate the requirement that both templates remain separate and reusable. This aligns with AZ-104 Deploy and manage Azure compute resources , specifically Automate deployment of resources by using Azure Resource Manager templates or Bicep files .

質問: **162**

Azureサブスクリプションがあります。  
ユーザーは、ホームまたは顧客サイトからサブスクリプションのリソースにアクセスします。ユーザーは自宅から、AzureリソースにアクセスするためにポイントツーサイトVPNを確立する必要があります。顧客サイトのユーザーは、サイト間VPNを使用してAzureリソースにアクセスします。  
複数のAzure仮想マシンで実行されるApp1という名前の基幹業務アプリがあります。仮想マシンはWindows Server 2016を実行します。  
App1への接続がすべての仮想マシンに分散していることを確認する必要があります。  
使用できる2つのAzureサービスは何ですか？それぞれの正解は完全なソリューションを示します。  
注 :それぞれの正しい選択は1ポイントの価値があります。

- A. パブリックロードバランサー
- B. トラフィックマネージャー
- C. Azureコンテンツ配信ネットワーク (CDN)
- D. 内部ロードバランサー
- E. Azureアプリケーションゲートウェイ

正解: ([正解を表示します](#))

Line of Business WebAPP works on VMs need internal load balancer. So D is needed. Then deploy WebAPP on VMs, check the link. <https://docs.microsoft.com/en-us/azure/application-gateway/quick-create-portal> So B is needed as well. The orignal answer is not accomplished.

質問: **163**

Subscription1 という名前の Azure サブスクリプションがあります。  
contosostorage という名前の Azure ストレージ アカウントを作成し、次に data という名前のファイル共有を作成します。  
データファイル共有のファイルを参照するスクリプトには、どのUNCパスを含めるべきですか？答えは、適切な値を正しいターゲットにドラッグすることです。各値は、1回、複数回、または全く使用されない場合があります。  
コンテンツを表示するには、ペイン間の分割バーをドラッグするか、スクロールする必要がある場合があります。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Values

blob

blob.core.windows.net

contosostorage

data

file

file.core.windows.net

portal.azure.com

subscription1

Answer Area

.

\

正解:

Values

blob

blob.core.windows.net

contosostorage

data

file

file.core.windows.net

portal.azure.com

subscription1

Answer Area

contosostorage

.

file.core.windows.net

\

data

Explanation:



Azure File Shares are accessed using a UNC (Universal Naming Convention) path, which follows a strict and well-defined format in Microsoft Azure. When you create an Azure Storage account and then create a file share within that account, the file share is exposed over the SMB protocol and can be mounted or referenced just like a traditional Windows file share.

According to Microsoft Azure Administrator documentation and study guides, the UNC path format for an Azure file share is:

\\ < storage-account-name > .file.core.windows.net\ < file-share-name > In this scenario:

- \* The Azure subscription name (Subscription1) is not part of the UNC path.
- \* The storage account name is contosostorage.
- \* The file share name is data.
- \* Azure Files always uses the file.core.windows.net endpoint for SMB-based file shares.
- \* Blob endpoints (blob.core.windows.net) are used only for Blob Storage and are not valid for file shares.

Putting these components together results in the correct UNC path that can be used inside scripts, PowerShell, batch files, or application configurations to reference files stored in the Azure file share.

質問: 164

次の表に示す仮想ネットワークを含む Azure サブスクリプションがあります。

| Name  | Peered with | DNS server               |
|-------|-------------|--------------------------|
| VNET1 | VNET2       | Default (Azure-provided) |
| VNET2 | VNET1       | 10.10.0.4                |

次の表に示す仮想マシンがあります。

| Name    | IP address  | Network interface | Connects to   |
|---------|-------------|-------------------|---------------|
| Server1 | 10.10.0.4   | NIC1              | VNET1/Subnet1 |
| Server2 | 172.16.0.4  | NIC2              | VNET1/Subnet2 |
| Server3 | 192.168.0.4 | NIC3              | VNET2/Subnet2 |

次の表に示す仮想ネットワーク インターフェイスがあります。

| Name | DNS server                   |
|------|------------------------------|
| NIC1 | Inherit from virtual network |
| NIC2 | 10.10.0.4                    |
| NIC3 | Inherit from virtual network |

Server1 は、次の表に示すリソースを含む DNS サーバーです。

| Name              | Type             | Value          |
|-------------------|------------------|----------------|
| contoso.com       | Primary DNS zone | Not applicable |
| Host1.contoso.com | A record         | 131.107.10.15  |

contoso.com という名前の Azure プライベート DNS ゾーンがあり、そこには VNET2 への仮想ネットワーク リンクと、次の表に示すレコードがあります。

| Name  | Type     | Value          |
|-------|----------|----------------|
| Host1 | A record | 131.107.200.20 |
| Host2 | A record | 131.107.50.50  |

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Statements

Microsoft

Yes

No

Server2 resolves host2.contoso.com to 131.107.50.50.

☐

☐

Server2 resolves host1.contoso.com to 131.107.10.15.

☐

☐

Server3 resolves host2.contoso.com to 131.107.50.50.

☐

☐

正解:

Statements

Microsoft

Yes

No

Server2 resolves host2.contoso.com to 131.107.50.50.

☐

☒

Server2 resolves host1.contoso.com to 131.107.10.15.

☒

☐

Server3 resolves host2.contoso.com to 131.107.50.50.

☐

☒

Explanation:

NYN

質問: 165

storage1 という Azure Storage アカウントがあり、このアカウントには container 1 と container2 という 2 つのコンテナが含まれています。両方のコンテナで BLOB のバージョン管理が有効になっています。重要な BLOB の BLOB スナップショットを定期的 to 取得します。次のライフサイクル管理ポリシーを作成します。

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "version": {
            "tierToCool": {
              "daysAfterCreationGreaterThan": 15
            },
            "tierToArchive": {
              "daysAfterLastTierChangeGreaterThan": 7,
              "daysAfterCreationGreaterThan": 30
            }
          }
        },
        "filters": {
          "blobTypes": [
            "blockBlob"
          ],
          "prefixMatch": {
            "container1/"
          }
        }
      }
    }
  ]
}
```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

| Statements                                                                                 | Yes                   | No                    |
|--------------------------------------------------------------------------------------------|-----------------------|-----------------------|
| A blob snapshot automatically moves to the Cool access tier after 15 days.                 | <input type="radio"/> | <input type="radio"/> |
| A blob version in container2 automatically moves to the Archive access tier after 30 days. | <input type="radio"/> | <input type="radio"/> |
| A rehydrated version automatically moves to the Archive access tier after 30 days.         | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                                                                                 | Yes                              | No                               |
|--------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|
| A blob snapshot automatically moves to the Cool access tier after 15 days.                 | <input checked="" type="radio"/> | <input type="radio"/>            |
| A blob version in container2 automatically moves to the Archive access tier after 30 days. | <input type="radio"/>            | <input checked="" type="radio"/> |
| A rehydrated version automatically moves to the Archive access tier after 30 days.         | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

Based on the lifecycle management policy you created and the information from the web search results, here are the answers to your statements:

A blob snapshot automatically moves to the Cool access tier after 15 days. = Yes A blob version in container2 automatically moves to the Archive access tier after 30 days. = No A rehydrated version automatically moves to the Archive access tier after 30 days. = No The lifecycle management policy you created has two rules: one for container1 and one for container2. The rule for container1 has an action that moves blob snapshots to the Cool access tier if they are older than 15 days. Therefore, a blob snapshot in container1 will automatically move to the Cool access tier after 15 days, regardless of the access tier of the base blob. The rule for container2 has an action that moves blob versions to the Archive access tier if they are older than 30 days and have a prefix match of "archive/". Therefore, a blob version in container2 will only automatically move to the Archive access tier after 30 days if its name starts with "archive/". Otherwise, it will remain in its current access tier. A rehydrated version is a blob version that was previously in the Archive access tier and was restored to an online access tier (Hot or Cool) by using the rehydrate priority option1. A rehydrated version does not automatically move to the Archive access tier after 30 days, unless there is a lifecycle management policy rule that explicitly specifies this action. In your case, neither of the rules applies to rehydrated versions, so they will stay in their online access tiers until you manually change them or delete them.

質問: 166

Azure 仮想マシンを 3 つの Azure リージョンにデプロイします。  
各リージョンには仮想ネットワークが含まれます。各仮想ネットワークには、フルメッシュ トポロジでピアリングされた複数のサブネットが含まれています。  
各サブネットには、ルールが定義されたネットワーク セキュリティ グループ (NSG) が含まれています。  
ユーザーは、あるリージョンの仮想マシンから別のリージョンの仮想マシンに接続するためにポート 33000 を使用できないと報告しています。  
問題を診断するために使用できる 2 つのオプションはどれですか? それぞれの正解は完全な解決策を示します。  
注: 正しく選択するたびに 1 ポイントの価値があります。

- A. Azure 仮想ネットワーク マネージャー
- B. IP フロー検証
- C. Azure Monitor ネットワーク インサイト
- D. 接続のトラブルシューティング
- E. 選択的なセキュリティ ルール

正解: B,D (コメントを发表する)

In Azure Network Watcher, there are several diagnostic tools used to troubleshoot network connectivity and security issues.

\* IP Flow Verify - This feature allows you to test whether a packet is allowed or denied to or from a virtual machine based on Network Security Group (NSG) rules. It helps identify issues with network security configurations such as incorrect ports or blocked traffic.

\* Connection Troubleshoot - This tool validates whether a connection between a source and destination (such as between two VMs) can be established. It simulates network connectivity paths and identifies where the connection might be failing.

In the given scenario, a user cannot connect using port 33000 between virtual machines in different regions.

The correct diagnostic steps would involve verifying whether the port is being blocked by an NSG (using IP Flow Verify) and testing end-to-end connectivity (using Connection Troubleshoot).

According to Azure Network Watcher documentation:

"Use IP Flow Verify to determine if a packet is allowed or denied by an NSG rule."

"Use Connection Troubleshoot to check direct or peered network connectivity between virtual machines and analyze network path issues." Both tools provide complementary diagnostics to determine whether network traffic is correctly permitted or blocked.

質問: 167

次の表に示すストレージ アカウントを含む Azure サブスクリプションがあります。

| Name     | Performance | Premium account type  |
|----------|-------------|-----------------------|
| storage1 | Standard    | <i>Not applicable</i> |
| storage2 | Premium     | Block blobs           |
| storage3 | Premium     | File shares           |
| storage4 | Premium     | Page blobs            |

ホットアクセス層のBLOBストレージを必要とするApp1というアプリをデプロイする必要があります。App1に使用できるストレージアカウントはどれですか？

- A. ストレージ1のみ
- B. ストレージ1とストレージ4のみ
- C. ストレージ2とストレージ4のみ
- D. ストレージ1、ストレージ2、ストレージ4のみ
- E. ストレージ1、ストレージ2、ストレージ3、ストレージ4

正解: [\(正解を表示します\)](#)

Azure Blob Storage Hot access tier is supported by general-purpose storage accounts and premium block blob or premium page blob accounts, but not by premium file share-only accounts.

From Microsoft Azure Storage documentation:

Standard (general-purpose v2) storage accounts support Hot, Cool, and Archive access tiers.

Premium Block Blob accounts support blob workloads optimized for high transaction rates and low latency and effectively function in the Hot tier.

Premium Page Blob accounts also support blob workloads suitable for hot data scenarios.

Premium File Share accounts are strictly for Azure Files and do not support Blob storage tiers.

Evaluation of the table:

storage1 (Standard) # # Supports Hot tier

storage2 (Premium - Block blobs) # # Supports blob workloads (Hot equivalent) storage3 (Premium - File shares) # # No blob support storage4 (Premium - Page blobs) # # Supports blob workloads Final Answer:

D). storage1, storage2, and storage4 only

質問: 168

storage 1 という名前のストレージ アカウントを含む Azure サブスクリプションがあります。storage1 アカウントには container1 という名前のコンテナが含まれています。

90 日後に、含まれる BLOB を最も低コストの層に自動的に移動する、storage1 のライフサイクル管理ルールを作成する必要があります。

ルールをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

## Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToArchive": {
              "enableAutoTierToHotFromCool": {
                "tierToArchive": {
                  "tierToCool": {
                    "daysAfterModificationGreaterThan": 90
                  }
                }
              }
            }
          }
        }
      }
    }
  ]
}
```

Microsoft

...  
"filters": {  
 "prefixMatch": [  
 "blobIndexMatch": [  
 "blobTypes": [  
 "prefixMatch": [  
 "container1/"

...  
}

正解:

## Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToArchive": {
              "enableAutoTierToHotFromCool": {
                "tierToArchive": {
                  "tierToCool": {
                    "daysAfterModificationGreaterThan": 90
                  }
                }
              }
            }
          }
        }
      }
    }
  ]
  ...
  "filters": {
    "prefixMatch": [
      "blobIndexMatch": [
        "blobTypes": [
          "prefixMatch": [
            "container1/"
          ]
        ]
      ]
    ]
  }
  ...
}
```

Explanation:



質問: 169

VM1 という名前の仮想マシンを含む Azure サブスクリプションをお持ちです。VM1 に対して Azure Site Recovery を構成する予定です。Recovery Services コンテナを作成し、VM1 のレプリケーションを有効にします。フェイルオーバーが発生した場合に、VM1 がターゲットの Azure リージョンで自動的に起動するようにする必要があります。どうすればよいですか？

- A. Azure Automation アカウントを作成します。
- B. Azure ロジック アプリを構成します。
- C. VM1 の更新管理を有効にします。
- D. 復旧計画を設定します。

正解: (正解を表示します)

質問: 170

Azure サブスクリプションをお持ちです。

次の表に示す Azure コンテナ インスタンスをデプロイする予定です。

| Name      | Operating system                                |
|-----------|-------------------------------------------------|
| Instance1 | Nano Server installation of Windows Server 2019 |
| Instance2 | Server Core installation of Windows Server 2019 |
| Instance3 | Linux                                           |
| Instance4 | Linux                                           |

コンテナ グループにデプロイできるインスタンスはどれですか？

- A. インスタンス1のみ
- B. インスタンス2のみ
- C. インスタンス1とインスタンス2のみ
- D. インスタンス3とインスタンス4のみ

正解: [\(正解を表示します\)](#)

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-container-groups> Multi- container groups currently support only Linux containers. For Windows containers, Azure Container Instances only supports deployment of a single container instance. While we are working to bring all features to Windows containers, you can find current platform differences in the service

質問: 171

VM1 という名前の仮想マシンと、User1 および User2 という名前の 2 人のユーザーを含む Azure サブスクリプションがあります。

User1 に VM1 の共同作成者ロールを割り当てます。

User1 が VM1 で実行できる 2 つのアクションはどれですか。それぞれの正解は完全なソリューションを示します。

- A. データ ディスクを追加します。
- B. 毎日のバックアップを設定します。
- C. User2 に VM1 の共同作成者ロールを割り当てます。
- D. VM1 のイメージを Azure コンピューティング ギャラリーにアップロードします。
- E. User2 に VM1 のリーダー ロールを割り当てます。

正解: [\(正解を表示します\)](#)

In Microsoft Azure, Role-Based Access Control (RBAC) defines what actions users can perform on resources. The Contributor role is one of the built-in Azure roles that provides extensive permissions but excludes access management rights.

According to the Microsoft Azure Administrator documentation under "Built-in roles for Azure resources", the Contributor role has the following definition:

"Grants full access to manage all Azure resources, including the ability to create and manage resources, but does not grant permission to assign roles or manage access rights." This means that a Contributor can perform any configuration or management task on a resource (such as starting, stopping, resizing, or modifying a virtual machine) but cannot grant or modify permissions for other users.

Let's analyze each option in this context:

A). Add a data disk - # Correct

A Contributor can attach, detach, or manage data disks for the VM.

These actions involve resource management operations under the Microsoft.Compute/virtualMachines namespace, which are permitted to the Contributor role.

Microsoft Documentation Reference (Compute Resource Provider Operations):

Contributors can "create, update, and delete virtual machines, disks, and configurations." B). Configure a daily backup - # Incorrect Configuring a daily backup requires permissions within the Recovery Services vault or Backup vault resource type (Microsoft.RecoveryServices/vaults/\*).

The Contributor role on the VM does not automatically grant access to vault-level resources.

Hence, User1 cannot configure or schedule a daily backup without additional permissions on the vault.

C). Assign User2 the Contributor role for VM1 - # Incorrect

Assigning or modifying RBAC roles requires the Owner or User Access Administrator role because it involves operations under the Microsoft.Authorization/roleAssignments/\* permission.

The Contributor role does not include these permissions.

Therefore, User1 cannot assign roles to others.

D). Upload an image of VM1 to an Azure compute gallery - # Correct

The Contributor can capture or generalize a VM and upload its image to an Azure Compute Gallery, as this operation falls under the VM and image management permissions (Microsoft.Compute/galleries/images/\*).

This action is considered a resource management operation and does not require ownership or RBAC rights beyond Contributor.

E). Assign User2 the Reader role for VM1 - # Incorrect

Similar to option C, assigning RBAC roles is a management operation restricted to users with Owner or User Access Administrator privileges. Contributors cannot modify or assign access control permissions.

質問: 172

次の表に示すグループを含む Microsoft Entra テナントがあります。

| Name   | Type          | Has an assigned license |
|--------|---------------|-------------------------|
| Group1 | Security      | Yes                     |
| Group2 | Security      | No                      |
| Group3 | Microsoft 365 | Yes                     |
| Group4 | Microsoft 365 | No                      |

テナントには次の表に示すユーザーが含まれます。

| Name  | Member of | Has a direct assigned license |
|-------|-----------|-------------------------------|
| User1 | None      | Yes                           |
| User2 | Group1    | No                            |
| User3 | Group4    | Yes                           |
| User4 | None      | No                            |

削除できるユーザーとグループはどれですか? 回答するには、回答領域で適切なオプションを選択してください。  
注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Users:

User4 only  
User4 only  
User1 and User4 only  
User2 and User4 only  
User1, User2, User3, and User4

Groups:

Group2 and Group4 only  
Group2 only  
Group2 and Group3 only  
Group2 and Group4 only  
Group1, Group2, Group3, and Group4

正解:  
Answer Area

Users:

User4 only  
User4 only  
User1 and User4 only  
User2 and User4 only  
User1, User2, User3, and User4

Groups:

Group2 and Group4 only  
Group2 only  
Group2 and Group3 only  
Group2 and Group4 only  
Group1, Group2, Group3, and Group4

Explanation:



質問: 173

財務部門のユーザー向け設定を自動化するソリューションを推奨する必要があります。ソリューションは技術要件を満たす必要があります。推奨には何を含めるべきでしょうか？

- A. Azure AP B2C
- B. Azure AD ID 保護
- C. Azure ロジック アプリと Microsoft Identity Management (MIM) クライアント
- D. 動的グループと条件付きアクセスポリシー

正解: D (コメントを发表する)

Technically, The finance department needs to migrate their users from AD to AAD using AADC based on the finance OU, and need to enforce MFA use. This is conditional access policy. Employees also often get promotions and/or join other departments and when that occurs, the user's OU attribute will change when the admin puts the user in a new OU, and the dynamic group conditional access exception (OU= [Department Name Value]) will move the user to the appropriate dynamic group on next AADC delta sync.

<https://docs.microsoft.com/en-us/azure/active-directory/enterprise-users/groups-dynamic-membership>

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/overview>

<https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-userstates>

質問: 174

VM1、VM2、VM3という3つの仮想マシンを含むAzureサブスクリプションがあります。これらの仮想マシンはすべて、AVSet1という可用性セットに格納されています。VM1を新しい仮想マシンサイズにスケールアップする必要がありますが、目的のサイズが利用できません。まず何をすべきでしょうか？

- A. VM2 と VM3 をシャットダウンします。
- B. AVSet1 を管理された可用性セットに変換します。
- C. VM1 の割り当てを解除します。
- D. 近接配置グループを作成します。

正解: C (コメントを发表する)

When resizing a virtual machine in Azure, the target size must be available on the same hardware cluster where the VM currently resides. If the desired size is unavailable, Azure prevents the resize operation.

According to the Microsoft Azure Administrator guide:

"To resize a virtual machine to a size not available on the current hardware cluster, the VM must first be deallocated. When a VM is deallocated, its current hardware association is released, allowing Azure to place it on a different cluster that supports the new size." In this scenario, VM1 is part of an availability set (AVSet1). All VMs in an availability set must reside on the same hardware cluster. Therefore, to resize VM1 successfully to a new size that isn't currently available, you must first deallocate the VM, which frees it from its existing cluster allocation.

After deallocation, you can resize it, and Azure will allocate it to a compatible cluster.

質問: 175

storage1 という名前の Azure Storage アカウントがあり、そこには BLOB コンテナが含まれています。BLOB コンテナのデフォルトのアクセス層は Hot です。Storage1 には container! という名前のコンテナが含まれています。

次の表に示すように、storage1 にライフサイクル管理ルールを作成します。

| Name  | Rule scope                      | Blob type   | Blob subtype | Rule block                                                                                                                                       | Prefix match    |
|-------|---------------------------------|-------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Rule1 | Limit blobs by using filters.   | Block blobs | Base blobs   | If base blobs were not modified for two days, move to archive storage.<br>If base blobs were not modified for nine days, delete the blob.        | container1/Dep1 |
| Rule2 | Apply to all blobs in storage1. | Block blobs | Base blobs   | If base blobs were not modified for three days, move to cool storage.<br>If base blobs were not modified for nine days, move to archive storage. | Not applicable  |

次の表に示すアクションを実行します。

| Date      | Action                                                                             |
|-----------|------------------------------------------------------------------------------------|
| October 1 | Upload three files named Dep1File1.docx, File2.docx, and File3.docx to container1. |
| October 2 | Edit Dep1File1.docx and File3.docx.                                                |
| October 5 | Edit File2.docx.                                                                   |

以下の各記述について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。注：正しい選択肢は1つにつき1点です。

Answer Area

| Statements                                                  | Yes                   | No                    |
|-------------------------------------------------------------|-----------------------|-----------------------|
| On October 10, you can read Dep1File1.docx without a delay. | <input type="radio"/> | <input type="radio"/> |
| On October 10, you can read File2.docx without a delay.     | <input type="radio"/> | <input type="radio"/> |
| On October 10, you can read File3.docx without a delay.     | <input type="radio"/> | <input type="radio"/> |

正解:

| Statements                                                  | Yes                              | No                               |
|-------------------------------------------------------------|----------------------------------|----------------------------------|
| On October 10, you can read Dep1File1.docx without a delay. | <input type="radio"/>            | <input checked="" type="radio"/> |
| On October 10, you can read File2.docx without a delay.     | <input checked="" type="radio"/> | <input type="radio"/>            |
| On October 10, you can read File3.docx without a delay.     | <input type="radio"/>            | <input checked="" type="radio"/> |

Explanation:

On October 10, you can read Dep1File1.docx. = NO

Dep1File1.docx is a blob in container1 that was uploaded on October 1 and edited on October 2. According to the lifecycle management rule 1, any blob in container1 that has not been modified for 7 days will be moved to the archive tier. Therefore, on October 9, Dep1File1.docx will be moved to the archive tier. Blobs in the archive tier cannot be read unless they are first rehydrated, which may take several hours or days. Therefore, on October 10, you cannot read Dep1File1.docx unless you rehydrate it first.

On October 10, you can read File2.docx. = YES

File2.docx is a blob in container1 that was uploaded on October 1 and edited on October 5. According to the lifecycle management rule 1, any blob in container1 that has not been modified for 7 days will be moved to the archive tier. Therefore, on October 12, File2.docx will be moved to the archive tier. However, on October 10, File2.docx is still in the hot tier, which means it can be read without any delay or cost.

On October 10, you can read File3.docx. = NO

File3.docx is a blob in container1 that was uploaded on October 1 and edited on October 2. According to the lifecycle management rule 2, any blob in container1 that has not been modified for 5 days will be deleted. Therefore, on October 7, File3.docx will be deleted from the storage account. Therefore, on October 10, you cannot read File3.docx because it no longer exists.

質問: 176

RSV1 という名前の Recovery Services コンテナがあります。RSV1 には、インスタント スナップショットを 5 日間保持し、毎日のバックアップを 14 日間保持するバックアップ ポリシーがあります。

RSV1 は VM1 のバックアップを毎日実行します。VM1 は、8 日前に更新された静的 Web サイトをホストしています。

VM1 を 8 日前の時点に回復する必要があります。ソリューションではダウンタイムを最小限に抑える必要があります。

まず何をすべきでしょうか？

- A. VM1 の割り当てを解除します。
- B. [既存の復元構成を置き換える] オプションを使用して VM1 を復元します。
- C. VM1 を削除します。
- D. [新しいリストア構成の作成] オプションを使用して VM1 をリストアします。

正解: [\(正解を表示します\)](#)

Azure Backup uses Recovery Services vaults (RSVs) to manage backup and restore operations for virtual machines. Each backup consists of:

Instant Restore Snapshots - retained for quick recovery (up to 5 days in this case).

Daily Recovery Points - stored in the Recovery Services vault based on the retention policy (14 days here).

In the scenario, VM1's website was updated eight days ago, meaning the restore point required is eight days old - beyond the 5-day instant snapshot retention period but within the 14-day daily backup retention window.

According to the Azure Backup documentation, restoring from a vault-stored daily recovery point (not instant snapshot) involves creating a new virtual machine because it ensures a non-disruptive restore and minimal downtime.

Here's how it works:

The backup is used to create a new VM in the same or different resource group.

Once the new VM is confirmed operational, you can redirect traffic or replace the original VM during a controlled maintenance window.

The "Replace existing" option, on the other hand, overwrites the current VM - leading to downtime and risk if validation fails.

Therefore, to restore VM1 to its state from eight days ago while minimizing downtime, you must first restore using the "Create new restore configuration" option, validate the restore, and then switch over seamlessly.

質問: 177

VM1 という名前の仮想マシンを含む Azure サブスクリプションがあります。

VM1 の CPU 使用率が 80% を超えたときにアラートをトリガーする Azure Monitor アラート ルールをデプロイする予定です。

アラート ルールによって、User1 と User2 という 2 人のユーザーに電子メール メッセージが送信されるようにする必要があります。

Azure Monitor 用に何を作成する必要がありますか？

- A. アクショングループ
- B. メールが有効なセキュリティ グループ
- C. 配布グループ
- D. Microsoft 365 グループ

正解: **A** ([コメントを發表する](#))

In Azure Monitor, an alert rule defines the condition (metric or log query) that triggers an alert, but notification and automation actions are managed through an Action Group.

An Action Group is a collection of notification preferences and actions used by Azure Monitor and Service Health alerts. It supports multiple actions such as:

- \* Email notifications
- \* SMS messages
- \* Push notifications
- \* Webhook triggers
- \* Automation runbooks

In this scenario, you must send an email message to two users (User1 and User2) whenever the CPU usage of VM1 exceeds 80%. To achieve this, you:

- \* Create an Action Group in Azure Monitor.
- \* Add two email receivers (User1 and User2) to that action group.
- \* Associate the Action Group with the alert rule monitoring VM1's CPU percentage.

Microsoft Official Documentation Extract (Paraphrased):

"Action groups are reusable notification collections that define how you want to be alerted when an alert rule triggers. You can specify multiple email recipients, SMS numbers, and webhook endpoints within one Action Group." (Source: Microsoft Learn - Create and manage action groups in Azure Monitor.)

質問: **178**

Windows Server 2019 を実行する VM1 という名前の Azure 仮想マシンがあります。

VM1 を Template1 という名前のテンプレートとして Azure Resource Manager ライブラリに保存します。

Template1 から VM2 という名前の仮想マシンを展開する予定です。

VM2 の展開中に何を構成できますか？

- A. 仮想マシンのサイズ
- B. オペレーティング システム
- C. 管理者のユーザー名
- D. リソース グループ

正解: ([正解を表示します](#))

Resource Group is the correct answer: Admin user, password, vm size and os are the part of ARM templates.

But resource group is not hence needs to be mentioned while deployment! Refer below sample ARM template for reference in which all above attributes passed in parameter. <https://github.com/Azure/azure-quickstart-templates/blob/master/101-vm-simple-windows/azuredeploy.json>

質問: **179**

Active Directory の問題を解決する必要があります。

何をすべきでしょうか？

- A. Active Directory ユーザーとコンピューターからユーザー アカウントを選択し、ユーザー プリンシパル名の値を変更します。
- B. idfix.exe を実行し、編集アクションを使用します。
- C. Active Directory ドメインと信頼関係から、UPN サフィックスのリストを変更します。

**D.** Azure AD Connect から、送信同期規則を変更します。

正解: [\(正解を表示します\)](#)

IdFix is used to perform discovery and remediation of identity objects and their attributes in an on-premises Active Directory environment in preparation for migration to Azure Active Directory. IdFix is intended for the Active Directory administrators responsible for directory synchronization with Azure Active Directory.

Scenario: Active Directory Issue

Several users in humongousinsurance.com have UPNs that contain special characters.

You suspect that some of the characters are unsupported in Azure AD.

References: <https://www.microsoft.com/en-us/download/details.aspx?id=36832>

質問: **180**

あなたは、webapp1 という名前の Web アプリを含む Azure サブスクリプションを持っています。webapp1 に www.contoso.com という名前のカスタム ドメインを追加する必要があります。最初に何をすべきですか？

**A.** 証明書をアップロードしてください。

**B.** 接続文字列を追加します。

**C.** ウェブアプリを停止します。

**D.** DNSレコードを作成します。

正解: **D** ([コメントを發表する](#))

When configuring a custom domain for an Azure Web App (App Service) such as webapp1, the first prerequisite is to create and configure a corresponding DNS record that maps your custom domain name (www.contoso.com) to the Azure App Service's default domain (webapp1.azurewebsites.net).

According to Microsoft Azure App Service documentation:

"Before you can assign a custom domain name to your web app, you must first create a CNAME record in your DNS provider that maps your custom domain (for example, www.contoso.com) to your app's default domain name (webappname.azurewebsites.net)." This DNS record ensures that requests to the custom domain are routed correctly to Azure's front-end load balancer hosting your web app. Only after the DNS record has propagated can you validate and bind the domain to the web app in the Azure portal.

Uploading an SSL certificate is only required for HTTPS bindings after the domain is added. Stopping the app or adding connection strings is unrelated to domain configuration.

Therefore, the first step is to create a DNS record at your domain registrar that points to your web app.

質問: **181**

storage1 という名前のストレージ アカウントを含む Azure サブスクリプションがあります。

ロールベースのアクセス制御 (RABC) ロールをストレージ 1 に割り当てるときに条件を使用する予定です。ロールを割り当てるときに条件をサポートするストレージ 1 サービスはどれですか。

**A.** コンテナのみ

**B.** ファイル共有のみ

**C.** テーブルのみ

**D.** キューのみ

**E.** コンテナとキューのみ

**F.** ファイル共有とテーブルのみ

正解: [\(正解を表示します\)](#)

Azure role-based access control (RBAC) now supports role assignment conditions for finer-grained access management. Conditions are written using the Azure Resource Manager (ARM) condition language and allow you to enforce specific rules (for example, limit access to particular blobs or queues).

However, conditional access in RBAC is currently available only for data actions in Azure Storage accounts and Azure Key Vault. According to the Microsoft Learn documentation for Azure Storage RBAC with conditions, the following services support conditional role assignments:

\* Blob storage (containers and blobs)

\* Queue storage

This means that you can apply conditions on containers (for blobs) and queues, but not on file shares or tables.

Conditions can restrict access to:

\* Specific container names or blob prefixes.

\* Specific queue names or messages.

For example, you could allow a user to read blobs only under a given prefix or queue, enhancing least- privilege control.

# Supported: Containers (Blob storage), Queues

# Not supported: File shares, Tables

Microsoft Azure Reference (Conceptual Summary):

"You can add conditions to Azure role assignments for blob and queue data actions. Conditions are not yet supported for Azure Files or Tables." (Source: Microsoft Learn - Azure role assignment conditions for storage data actions)

有効的な**AZ-104J**問題集はJPNTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 182

Sub1 という名前の Azure サブスクリプションがあり、その中に RG1 という名前のリソース グループがあります。RG1 内の現在および将来のすべてのリソースに Department =HR というタグが適用されていることを確認する必要があります。どうすればよいですか？

A. Azure ポリシーを作成して割り当てます。

B. RG1 の各リソースに Department =HR タグを適用します。

C. RG1 に Departments タグを直接適用します。

D. Sub1 に Departments タグを適用します。

正解: ([正解を表示します](#))

質問: 183

スコープ 1 の計画された変更を実装します。

スコープ 1 が技術要件を満たしていることを確認する必要があります。

Scope1 を使用して何を暗号化できますか？

A. ストレージ2内のコンテナとBLOBのみ

B. ストレージ1とストレージ2内のコンテナとBLOB

C. ストレージ2 内のコンテナ、BLOB、ファイル共有のみ

D. ストレージ1とストレージ2内のコンテナ、BLOB、ファイル共有

E. ストレージ2 内のコンテナ、BLOB、ファイル共有、キュー、テーブルのみ

正解: E ([コメントを發表する](#))

In Microsoft Azure, encryption scopes are a StorageV2 (general-purpose v2) storage account feature that allows fine-grained control over encryption settings for data stored within a single account. According to Microsoft Azure Storage documentation, an encryption scope defines a specific encryption context that can be applied at the container or blob level and is supported in non-hierarchical namespace storage accounts (those without Data Lake Gen2 enabled).

In the given scenario:

- \* storage1 has Hierarchical namespace = Yes (Data Lake Storage Gen2 enabled).
- \* storage2 has Hierarchical namespace = No.
- \* The plan was to create an encryption scope named Scope1 in storage2.
- \* The technical requirement specifies that Scope1 must be used to encrypt storage services.

According to the Azure Administrator documentation on encryption scopes:

"Encryption scopes are supported for block blobs, append blobs, page blobs, Azure Files, queues, and tables in standard StorageV2 accounts. Encryption scopes are not supported in hierarchical namespace (Data Lake Gen2) enabled accounts." This means that Scope1-created in storage2, which does not have hierarchical namespace-can encrypt all blob data (containers and blobs) as well as file shares, queues, and tables.

However, storage1 cannot use encryption scopes because hierarchical namespace storage accounts (ADLS Gen2) manage encryption at the account level and do not support per-scope encryption.

Therefore, only storage2 can apply Scope1, and it can encrypt containers, blobs, file shares, queues, and tables.

## Topic 2, Contoso LtdOverview

Contoso, Ltd. is a manufacturing company that has offices worldwide. Contoso works with partner organizations to bring products to market.

Contoso products are manufactured by using blueprint files that the company authors and maintains.

### Existing Environment

Currently, Contoso uses multiple types of servers for business operations, including the following:

File servers

Domain controllers

Microsoft SQL Server servers

Your network contains an Active Directory forest named contoso.com. All servers and client computers are joined to Active Directory.

You have a public-facing application named App1. App1 is comprised of the following three tiers:

A SQL database

A web front end

A processing middle tier

Each tier is comprised of five virtual machines. Users access the web front end by using HTTPS only.

### Requirements

#### Planned Changes

Contoso plans to implement the following changes to the infrastructure:

Move all the tiers of App1 to Azure.

Move the existing product blueprint files to Azure Blob storage.

Create a hybrid directory to support an upcoming Microsoft Office 365 migration project.

### Technical Requirements

Contoso must meet the following technical requirements:

Move all the virtual machines for App1 to Azure.

Minimize the number of open ports between the App1 tiers.

Ensure that all the virtual machines for App1 are protected by backups.

Copy the blueprint files to Azure over the Internet.

Ensure that the blueprint files are stored in the archive storage tier.

Ensure that partner access to the blueprint files is secured and temporary.

Prevent user passwords or hashes of passwords from being stored in Azure.

Use unmanaged standard storage for the hard disks of the virtual machines.

Ensure that when users join devices to Azure Active Directory (Azure AD), the users use a mobile phone to verify their identity.

Minimize administrative effort whenever possible.

User Requirements

Contoso identifies the following requirements for users:

Ensure that only users who are part of a group named Pilot can join devices to Azure AD.

Designate a new user named Admin1 as the service administrator of the Azure subscription.

Admin1 must receive email alerts regarding service outages.

Ensure that a new user named User3 can create network objects for the Azure subscription.

質問: 184

次の表のリソースを含む Azure サブスクリプションがあります。

| Name    | Type            | Details               |
|---------|-----------------|-----------------------|
| VNet1   | Virtual network | <i>Not applicable</i> |
| Subnet1 | Subnet          | Hosted on VNet1       |
| VM1     | Virtual machine | On Subnet1            |
| VM2     | Virtual machine | On Subnet1            |

VM1とVM2は同じテンプレートから展開され、リモートデスクトップを使用してアクセスする基幹業務アプリケーションをホストしています。図に示されているネットワークセキュリティグループ (NSG)を構成します。(図のボタンをクリックしてください。)

→ Move  Delete

Resource group ([change](#))  
**ProductionRG**

Location  
**North Europe**

Subscription ([change](#))  
**Production subscription**

Subscription ID  
**14d26092-8e42-4ea7-b770-9dcef70fb1ea**

Tags ([change](#))  
[Click here to add tags](#)

Security rules  
1 inbound, 1 outbound

Associated with  
0 subnets, 0 network interfaces



### Inbound security rules

| PRIORITY | NAME                          | PORT | PROTOCOL | SOURCE            | DESTINATION    | ACTION                                                                                          |
|----------|-------------------------------|------|----------|-------------------|----------------|-------------------------------------------------------------------------------------------------|
| 1500     | Port_80                       | 80   | TCP      | Internet          | Any            |  Deny ...   |
| 65000    | AllowVnetInBound              | Any  | Any      | VirtualNetwork    | VirtualNetwork |  Allow ... |
| 65001    | AllowAzureLoadBalancerInBound | Any  | Any      | AzureLoadBalancer | Any            |  Allow ... |
| 65500    | DenyAllBound                  | Any  | Any      | Any               | Any            |  Deny ...  |

### Outbound security rules

| PRIORITY | NAME                  | PORT | PROTOCOL | SOURCE         | DESTINATION    | ACTION                                                                                          |
|----------|-----------------------|------|----------|----------------|----------------|-------------------------------------------------------------------------------------------------|
| 1000     | DenyWebSites          | 80   | TCP      | Any            | Internet       |  Deny ...  |
| 65000    | AllowVnetOutBound     | Any  | Any      | VirtualNetwork | VirtualNetwork |  Allow ... |
| 65001    | AllowInternetOutBound | Any  | Any      | Any            | Internet       |  Allow ... |
| 65500    | DenyAllOutBound       | Any  | Any      | Any            | Any            |  Deny ...  |

VM1 および VM2 のユーザーがインターネット上の Web サイトにアクセスできないようにする必要があります。  
何をすべきでしょうか？

- A. NSG を Subnet1 に関連付けます。
- B. NSG とネットワーク インターフェースの関連付けを解除します。
- C. DenyWebSites 送信セキュリティ ルールを変更します。
- D. Port\_80 の受信セキュリティルールを変更する

正解: (正解を表示します)

Outbound rule "DenyWebSites" is setup correctly to block outbound internet traffic over port 80. In the screenshot it states, " Associated with: 0 subnets, 0 NIC ' s " , so you need to associate the NSG to Subnet1. You can associate or dissociate a network security group from a NIC or Subnet. Reference: <https://docs.microsoft.com/en-us/azure/virtual-network/manage-network-security-group>

質問: 185

VNet1という仮想ネットワークを含むSubscription1というAzureサブスクリプションがあります。次の表のユーザーを追加します。

| User  | Role                |
|-------|---------------------|
| User1 | Owner               |
| User2 | Security Admin      |
| User3 | Network Contributor |

各構成を実行できるのはどのユーザーですか？ 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Add a subnet to VNet1:

User1 and User3 only

User1 only

User3 only

User1 and User3 only

User2 and User3 only

User1, User2, and User3

Assign a user the Reader role to VNet1:

User1 only

User1 only

User2 only

User3 only

User1 and User2 only

User2 and User3 only

User1, User2, and User3

正解:

**Answer Area**

Add a subnet to VNet1:

Assign a user the Reader role to VNet1:

 Microsoft

Explanation:

**Answer Area**

Add a subnet to VNet1:

Assign a user the Reader role to VNet1:

 Microsoft

This question evaluates understanding of Azure Role-Based Access Control (RBAC) permissions for managing virtual networks (VNETs) and assigning Azure roles.

1. Understanding Each User's Role

User

Role

Key Capabilities

User1

Owner

Full access to all resources, including the ability to delegate access by assigning roles.

User2

Security Admin

Can manage security policies and settings (Azure Defender, Security Center) but cannot manage or configure networking resources.

User3

Network Contributor

Can manage networking resources, including VNets, subnets, and network interfaces, but cannot grant access or assign roles.

2. Task Analysis

Task 1: Add a subnet to VNet1

To add a subnet to a virtual network, the user must have permission to modify network resources.

The relevant Azure RBAC action is:

Microsoft.Network/virtualNetworks/subnets/write

Roles that include this permission:

Owner - has all permissions.

Network Contributor - can manage all network settings, including adding/removing subnets.

# Therefore, User1 (Owner) and User3 (Network Contributor) can perform this task.

# User2 (Security Admin) cannot, because that role does not include network configuration permissions.

The Answer: User1 and User3 only

Task 2: Assign a user the Reader role to VNet1

To assign roles or manage access permissions, the user must have RBAC management privileges, specifically:

Microsoft.Authorization/roleAssignments/write

Only users with roles that include role assignment permissions can do this - such as:

Owner

User Access Administrator

The Network Contributor and Security Admin roles do not have this privilege.

# Therefore, User1 (Owner) can assign the Reader role.

# User2 and User3 cannot assign RBAC permissions.

The Answer: User1 only

Microsoft Official Documentation Extracts (AZ-104 Study Guide and Azure Docs):

"The Owner role has full access to all resources, including the right to delegate access to others."

"The Network Contributor role allows management of networking resources but not access control assignments."

"The Security Admin role manages security policies and alerts but cannot create, modify, or delete network resources." (Source: Microsoft Learn # Azure built-in roles for RBAC)

# Final Verified Answers Summary

Configuration

Correct Answer

Explanation:

Add a subnet to VNet1

User1 and User3 only

Owner and Network Contributor can modify virtual network configurations.

Assign a user the Reader role to VNet1

User1 only

Only Owner can assign roles (has roleAssignments/write permission).

Therefore, the final verified answer is:

# Add a subnet to VNet1 # User1 and User3 only

# Assign Reader role to VNet1 # User1 only

質問: 186

Traffic Analytics が構成された Azure サブスクリプションがあります。  
次の設定を持つ VM1 という名前の新しい仮想マシンをデプロイします。

\* 地域 - 米国東部

\* 仮想ネットワーク: VNet1

\* NICネットワークセキュリティグループ: NSG1

Traffic Analytics を使用して VM1 トラフィックを監視する必要があります。  
どのような設定を構成する必要がありますか？

A. VM1の診断設定

B. VM1 の分析情報

C. NSG1のNSGフローログ

D. NSG1の診断設定

正解: ([正解を表示します](#))

Traffic Analytics analyzes the network security group (NSG) flow logs to provide insights into traffic flow in your Azure cloud1. NSG flow logs are a feature of Network Watcher that allows you to view information about ingress and egress IP traffic through an NSG2. To use Traffic Analytics, you need to enable NSG flow logs for the network security groups you want to monitor1.

Diagnostic settings for VM1 or NSG1 are not required for Traffic Analytics. Diagnostic settings are used to stream log data from an Azure resource to different destinations such as Log Analytics workspace, Event Hubs, or Storage account3. Insights for VM1 are also not required for Traffic Analytics. Insights are a feature of Azure Monitor that provide analysis of the performance and health of an Azure resource4.

質問: 187

2 つの実行中のインスタンスを含む App1 という名前の Azure App Service アプリがあります。  
次の図に示すように自動スケール ルールが構成されています。

## Criteria

Metric namespace \*

Standard metrics

Metric name

Memory Percentage

1 minute time grain

Dimension Name

Operator

Dimension Values

Add

Instance

=

▼

All values

▼

+

If you select multiple values for a dimension, autoscale will aggregate the metric across the selected values, not evaluate the metric for each values individually.



MemoryPercentage (Average)

39.28 %

☐ Enable metric divide by instance count ⓘ

Operator \*

Greater than

▼

Metric threshold to trigger scale action \* ⓘ

70

%

Duration (minutes) \* ⓘ

15

✓

Time grain (minutes) ⓘ

1

Time grain statistic \* ⓘ

Average

▼

Time aggregation \* ⓘ

Average

▼

## Action


Operation \*

Increase count by

▼

Cool down (minutes) \* ⓘ

5



increase count by

instance count \*

1

[インスタンス制限のスケール条件] 設定では、[最大] を 5 に設定します。

30 分間に、App1 は使用可能なメモリの 80% を使用します。

30 分間の App1 のインスタンスの最大数はどれくらいですか？

- A. 2
- B. 3
- C. 4
- D. 5

正解: ([正解を表示します](#))

In Azure App Service, autoscale rules automatically adjust the number of running instances based on performance metrics such as CPU, memory, or custom metrics. The configuration shown uses the Memory Percentage metric with a threshold of 70%, meaning that if average memory utilization exceeds 70% for 15 minutes, Azure will trigger a scale-out action.

From the exhibit, the rule specifies:

- \* Metric threshold: Greater than 70%
- \* Duration: 15 minutes
- \* Action: Increase count by 1 instance
- \* Cool down: 5 minutes
- \* Maximum instance limit: 5

App1 currently has 2 running instances.

If App1 maintains 80% memory utilization for 30 minutes, the autoscale mechanism will trigger the scale-out action after each qualifying 15-minute window, adding 1 instance every 5-minute cooldown period until the maximum instance limit (5) is reached.

Therefore, the progression would be:

- \* Start: 2 instances
- \* After first 15 minutes: +1 instance # 3 total
- \* After next 5-minute cooldown and continuing high memory # +1 instance # 4 total
- \* After another cycle # +1 instance # 5 total

Since the rule continues to trigger until the maximum (5) is reached, and the memory usage remains above the threshold, the maximum number of instances that App1 will scale to during the 30-minute period is 5.

質問: 188

Azure サブスクリプションをお持ちです。

Templatejson という名前の次の Azure Resource Manager (ARM) テンプレートを作成します。

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {},
  "variables": {},
  "resources": [
    {
      "type": "Microsoft.Resources/resourceGroups",
      "apiVersion": "2022-12-01",
      "location": "eastus",
      "name": "Marketing"
    }
  ],
  "outputs": {}
}
```

Template.json をデプロイする必要があります。

Azure Cloud Shell からどの PowerShell コマンドレットを実行する必要がありますか？

A. 新しい AzManagementGroupDeployment

B. 新しい AzTenantDeployment

C. 新しい AzResourceGroupDeployment

D. 新しい AzSubscriptionDeployment

正解: (正解を表示します)

The provided ARM template creates a resource of type Microsoft.Resources/resourceGroups. Templates that create or manage resource groups must be deployed at the subscription scope, not at the resource group scope.

Microsoft guidance on deployment scopes explains that resource group creation requires a subscription-level deployment.

Accordingly, the correct Azure PowerShell cmdlet to deploy such a template from Azure Cloud Shell is New- AzSubscriptionDeployment. The other options target tenant, management group, or resource group scopes and are not appropriate for creating a resource group.

有効的な**AZ-104J**問題集はJPNTTest.com提供され、**AZ-104J**試験に合格することに役に立ちます！JPNTTest.comは今最新**AZ-104J**試験問題集を提供します。JPNTTest.com AZ-104J試験問題集はもう更新されました。ここで**AZ-104J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/AZ-104J-mondaishu> 481問、30%ディスカウント、特別な割引コード: **JPNshiken**」